



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM
VIRTUALIZED SERVICE ROUTER**

**INTERFACE CONFIGURATION GUIDE
RELEASE 16.0.R4**

3HE 14128 AAAB TQZZA 01

Issue: 01

November 2018

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or tradenames of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2018 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1	Getting Started	9
1.1	About This Guide	9
1.2	Interface Configuration Process	11
2	Interfaces	13
2.1	Configuration Overview	13
2.1.1	Chassis Slots and Cards	13
2.1.2	MCMs	15
2.1.3	MDA-a, MDA-aXP, MDA, MDA-XP and MDA-e Modules.....	15
2.1.4	XMAs/C-XMAs.....	18
2.1.5	Hardware Licensing.....	19
2.1.6	Software License Activation	21
2.1.7	Software License Records.....	23
2.1.8	CMAs.....	24
2.1.9	Versatile Service Module (VSM).....	25
2.1.10	Oversubscribed Ethernet MDAs	25
2.1.10.1	Rate Limiting.....	25
2.1.10.2	Packet Classification and Scheduling.....	26
2.1.11	Channelized MDA/CMA Support.....	28
2.1.11.1	Channelized DS-1/E-1 CMA.....	28
2.1.11.2	Channelized DS-3/E-3 CMA.....	28
2.1.11.3	Channelized Any Service Any Port (ASAP) CHOC-3/STM-1	28
2.1.11.4	Channelized OC-12/STM-4 ASAP MDAs.....	29
2.1.11.5	Channelized DS-3/E-3 ASAP MDA (4-Port).....	29
2.1.11.6	Channelized DS-3/E-3 ASAP MDA (12-Port).....	29
2.1.11.7	Channelized OC-3/STM-1 Circuit Emulation Services (CES) CMA and MDA.....	30
2.1.11.8	Network Interconnections	31
2.2	Digital Diagnostics Monitoring	32
2.2.1	SFPs and XFPs.....	36
2.2.2	Statistics Collection	36
2.3	Ports	37
2.3.1	Port Types	37
2.3.2	Port Features.....	40
2.3.2.1	Port State and Operational State.....	40
2.3.2.2	802.1x Network Access Control	41
2.3.2.3	MACsec.....	46
2.3.2.4	SONET/SDH Port Attributes.....	66
2.3.2.5	SONET/SDH Path Attributes	67
2.3.2.6	Multilink Frame Relay	68
2.3.2.7	FRF.12 End-to-End Fragmentation	71
2.3.2.8	FRF.12 UNI/NNI Link Fragmentation	72
2.3.2.9	MLFR/FRF.12 Support of APS, BFD, and Mirroring Features	73
2.3.2.10	Multilink Point-to-Point Protocol (MLPPP)	74
2.3.2.11	Multi-Class MLPPP.....	78

2.3.2.12	Cisco HDLC	85
2.3.2.13	Automatic Protection Switching (APS)	87
2.3.2.14	Inverse Multiplexing Over ATM (IMA)	116
2.3.2.15	Ethernet Local Management Interface (E-LMI)	119
2.3.2.16	Link Layer Discovery Protocol (LLDP)	120
2.3.2.17	Exponential Port Dampening	124
2.3.3	Per Port Aggregate Egress Queue Statistics Monitoring	127
2.4	FP4 Data Path Mapping	130
2.5	Port Cross-Connect (PXC)	131
2.5.1	PXC Terminology	132
2.5.2	PXC - Physical Port in Cross-Connect (Loopback) Mode	133
2.5.2.1	Operational State	134
2.5.3	PXC Sub-Ports	134
2.5.3.1	PXC Sub-Port Operational State	136
2.5.4	Port Statistics	137
2.5.4.1	Statistics on Physical PXC Ports	137
2.5.5	LAG with PXC Ports – PXC LAG	138
2.5.6	Basic PXC Provisioning	139
2.5.7	QoS	140
2.5.7.1	Queue Allocation on PXC Sub-Ports	140
2.5.7.2	Pool Allocations on PXC Ports	141
2.5.7.3	QoS Summary	141
2.5.8	Mirroring and LI on PXC Ports	142
2.5.9	Multi-Chassis Redundancy	142
2.5.10	Health Monitoring on the PXC Sub-Ports	142
2.5.11	Configuration Example	143
2.6	Forwarding Path Extensions (FPE)	146
2.7	LAG	148
2.7.1	LACP	148
2.7.1.1	LACP Multiplexing	149
2.7.1.2	LACP Tunneling	149
2.7.2	Active-Standby LAG Operation	150
2.7.3	LAG on Access QoS Consideration	152
2.7.3.1	Adapt QoS Modes	152
2.7.3.2	Per-fp-ing-queuing	154
2.7.3.3	Per-fp-egr-queuing	154
2.7.3.4	Per-fp-sap-instance	155
2.7.4	LAG and ECMP Hashing	156
2.7.4.1	Per Flow Hashing	157
2.7.4.2	Per Link Hashing	164
2.7.4.3	Explicit Per Link Hash Using LAG Link Mapping Profiles	166
2.7.4.4	Consistent Per Service Hashing	168
2.7.4.5	ESM – LAG Hashing per Vport	169
2.7.5	LAG Hold Down Timers	172
2.7.6	BFD over LAG Links	173
2.7.7	Mixed Port-Speed LAG Support	173
2.7.7.1	LAG Upgrade	175
2.7.8	Multi-Chassis LAG	176
2.7.8.1	Overview	176

2.7.8.2	MC-LAG and Subscriber Routed Redundancy Protocol (SRRP)	180
2.7.8.3	Point-to-Point (p2p) Redundant Connection Across Layer 2/3 VPN Network	180
2.7.8.4	DSLAM Dual Homing in Layer 2/3 TPSDA Model	182
2.8	G.8031 Protected Ethernet Tunnels	183
2.9	G.8032 Protected Ethernet Rings	184
2.10	Ethernet Port Monitoring	185
2.11	802.3ah OAM	189
2.11.1	OAM Events	191
2.11.1.1	Link Monitoring	192
2.11.2	Remote Loopback	198
2.11.3	802.3ah OAM PDU Tunneling for Epipe Service	199
2.11.3.1	802.3ah Grace Announcement	199
2.12	MTU Configuration Guidelines	206
2.12.1	Default MTU Values	206
2.12.2	Modifying MTU Defaults	207
2.12.3	Configuration Example	207
2.13	Deploying Preprovisioned Components	209
2.14	Setting Fabric Speed	210
2.14.1	fabric-speed-a (SFM5/CPM5)	210
2.14.2	fabric-speed-b (SFM5/CPM5)	210
2.14.3	fabric-speed-c (SFM5/CPM5)	210
2.15	Versatile Service Module	211
2.15.1	VSM-CCA-XP	211
2.16	Configuration Process Overview	214
2.17	Configuration Notes	215
2.18	Configuring Physical Ports with CLI	217
2.18.1	Preprovisioning Guidelines	217
2.18.1.1	Predefining Entities	217
2.18.1.2	Preprovisioning a Port	217
2.18.1.3	Maximizing Bandwidth Use	218
2.18.2	Basic Configuration	219
2.18.3	Common Configuration Tasks	221
2.18.3.1	Configuring Cards and MDAs	221
2.18.3.2	Configuring Cards, MCMs, and MCAs	223
2.18.3.3	Configuring Ports	224
2.19	Service Management Tasks	281
2.19.1	Modifying or Deleting an MDA, MCM, CMA or XMA	281
2.19.2	Modifying a Card Type	281
2.19.3	Deleting a Card	282
2.19.4	Deleting Port Parameters	283
2.19.5	Soft IOM Reset	283
2.19.5.1	Soft Reset	283
2.19.5.2	Deferred MDA Reset	284
2.20	Configuration Command Reference	287
2.20.1	Command Hierarchies	287
2.20.1.1	Card Commands	287
2.20.1.2	MCM Commands	288
2.20.1.3	MDA Commands	288

2.20.1.4	Power Commands.....	290
2.20.1.5	Virtual Scheduler Commands.....	290
2.20.1.6	Forwarding Plane (FP) Commands.....	290
2.20.1.7	Port Configuration Commands.....	292
2.20.1.8	Port XC Commands.....	295
2.20.1.9	Forwarding Path Extension (FPE) Commands.....	295
2.20.1.10	Port APS Commands.....	296
2.20.1.11	Ethernet Commands.....	296
2.20.1.12	HSQ Port Configuration Commands.....	306
2.20.1.13	Interface Group Handler Commands.....	306
2.20.1.14	Multilink Bundle Commands.....	307
2.20.1.15	SONET-SDH Commands.....	308
2.20.1.16	TDM Commands.....	311
2.20.1.17	DS3 Commands.....	313
2.20.1.18	E1 Commands.....	315
2.20.1.19	E3 Commands.....	317
2.20.1.20	LAG Commands.....	318
2.20.1.21	MACsec Commands.....	321
2.20.1.22	Ethernet Tunnel Commands.....	322
2.20.1.23	Multi-Chassis Redundancy Commands.....	323
2.20.2	Configuration Command Descriptions.....	325
2.20.2.1	Generic Commands.....	326
2.20.2.2	Card Commands.....	328
2.20.2.3	MCM Commands.....	332
2.20.2.4	MDA (XMA) Commands.....	333
2.20.2.5	MDA/Port QoS Commands.....	340
2.20.2.6	Power Commands.....	346
2.20.2.7	Virtual Scheduler Commands.....	349
2.20.2.8	Forwarding Plane Configuration Commands.....	353
2.20.2.9	MACsec Commands.....	379
2.20.2.10	General Port Commands.....	388
2.20.2.11	Port XC Commands.....	433
2.20.2.12	Forwarding Path Extension (FPE) Commands.....	435
2.20.2.13	APS Commands.....	439
2.20.2.14	Ethernet Port Commands.....	446
2.20.2.15	802.1x Port Commands.....	512
2.20.2.16	LLDP Port Commands.....	523
2.20.2.17	Network Port Commands.....	526
2.20.2.18	HSQ Port Commands.....	529
2.20.2.19	Interface Group Handler Commands.....	534
2.20.2.20	Multilink-Bundle Port Commands.....	536
2.20.2.21	SONET/SDH Port Commands.....	553
2.20.2.22	SONET/SDH Path Commands.....	560
2.20.2.23	ATM Interface Commands.....	568
2.20.2.24	Frame Relay Commands.....	572
2.20.2.25	TDM Commands.....	579
2.20.2.26	LAG Commands.....	600
2.20.2.27	Eth Tunnel Commands.....	620
2.20.2.28	ETH-CFM Configuration Commands.....	631

2.20.2.29	Multi-Chassis Redundancy Commands	645
2.20.2.30	Forwarding Plane Tools Commands	665
2.21	Show, Monitor, Clear, Debug, and Tools Command Reference	667
2.21.1	Command Hierarchies	667
2.21.1.1	Show Commands	667
2.21.1.2	Monitor Commands	670
2.21.1.3	Clear Commands.....	670
2.21.1.4	Debug Commands.....	671
2.21.1.5	Tools Commands	671
2.21.2	Command Descriptions	673
2.21.2.1	Hardware Show Commands.....	673
2.21.2.2	PEQ Show Commands.....	743
2.21.2.3	APS Show Commands	751
2.21.2.4	Port Show Commands.....	756
2.21.2.5	Multilink Bundle Show Commands	854
2.21.2.6	LAG Show Commands	875
2.21.2.7	Licensing Show Command.....	886
2.21.2.8	MACsec Show Commands.....	888
2.21.2.9	Monitor Commands	899
2.21.2.10	Clear Commands.....	908
2.21.2.11	Tools Commands	913
3	Standards and Protocol Support	939

1 Getting Started

1.1 About This Guide

This guide describes system concepts and provides configuration examples to provision Input/Output modules (IOMs), XMA Control Modules (XCMs), also referred to as cards, Media Dependent Adapters (MDAs), XRS Media Adapters (XMAs), and ports.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e • 7750 SR-7s/14s • 7750 SR-1	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the SR OS 16.0.Rx Software Release Notes, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.



Note: This guide generically covers Release 16.0.Rx content and may contain some content that will be released in later maintenance loads. Refer to the SR OS 16.0.Rx Software Release Notes, part number 3HE 14220 000x TQZZA or the *VSR Release Notes*, part number 3HE 14204 000x TQZZA, for information about features supported in each load of the Release 16.0.Rx software.

1.2 Interface Configuration Process

[Table 2](#) lists the tasks necessary to configure IOMs and XCMs (also referred to as cards), MDAs and XMAAs, and ports.



Note: For consistency across platforms, XMAAs are modeled in the SR OS (CLI and SNMP) as MDAs.

Unless specified otherwise:

- the term “card” is used generically to refer to both IOMs and XCMs
- the term “MDA” is used generically to refer to both MDAs and XMAAs

Table 2 Configuration Process

Area	Task	Section
Provisioning	Chassis slots and cards	Chassis Slots and Cards
	MCMs	MCMs
	MDAs	MDA-a, MDA-aXP, MDA, MDA-XP and MDA-e Modules
	Versatile Service Module	Versatile Service Module (VSM)
	Ports	Ports
Interface Configuration	MTU Configuration	MTU Configuration Guidelines
	Configure fabric speed	Setting Fabric Speed
	Preprovisioning	Preprovisioning Guidelines
	Configure cards and MDAs	Configuring Cards and MDAs
	Configure cards, MCMs, and MDAs	Configuring Cards, MCMs, and MCAs
	Configure ports	Configuring Ports
	Service management	Service Management Tasks

2 Interfaces

2.1 Configuration Overview



Note: This document uses the term “preprovisioning” in the context of preparing or preconfiguring entities such as chassis slots, cards, Media Dependent Adapters (MDAs), compact media adapters (CMAs), ports, and interfaces, prior to initialization. These entities can be installed while remaining administratively disabled (shutdown). When the entity is in a no shutdown state (administratively enabled), then the entity is considered to be provisioned.



Note: For consistency across platforms, XRS Media Adapters (XMAs) and Compact XMAs (C-XMAs) are modeled as MDAs.

Unless specified otherwise:

- the term “card” is used generically to refer to both Input Output Modules (IOMs) and XCMs
- the term “MDA” is used generically to refer to both MDAs and XMAs

Nokia routers provide the capability to configure chassis slots to accept specific card and MDA types and set the relevant configurations before the equipment is actually installed. The preprovisioning capability allows you to plan your configurations as well as monitor and manage your router hardware inventory. Ports and interfaces can also be preprovisioned. When the functionality is needed, the cards can be inserted into the appropriate chassis slots when required.

2.1.1 Chassis Slots and Cards

To preprovision a chassis slot, the card type must be specified. Operators can enter card type information for each slot. When a card is installed in a slot and enabled, the system verifies that the installed card type matches the provisioned card type. If the parameters do not match, the card remains offline. A preprovisioned slot can remain empty without conflicting with populated slots.

The general syntax for the configuration of card slots is similar for all platforms, though the number of available slots varies by platform and chassis model. The supported card-types vary by chassis. Refer to the appropriate platform Installation Guide for more information.

The 7950 XRS platforms accept XCMs in slots. An XCM has two slots, each of which accept an XMA or C-XMA module. The C-XMA modules require a mechanical adapter to fit in an XMA slot.

In the config context, use the following CLI commands and syntax examples to provision the chassis slot and XCM:

```
A:XRS20>config# card 1
A:XRS20>config>card# card-type xcm-x20
```

The 7450 ESS-7/12, and 7750 SR-7/12, and 7750 SR-12e platforms support a variety of IOM types (including the IOM3-XP, IOM3-XP-B, IOM3-XP-C, IOM4-e and IOM4-e-B) in designated chassis slots. IOMs have two slots for pluggable MDAs. The IOM3-XP, IOM3-XP-B and IOM3-XP-C support MDA and MDA-XP. The IOM4-e and IOM4-e-B support MDA-e modules.

In the config context, use the following CLI commands and syntax examples to provision a chassis slot and an IOM:

```
A:SR12-1>config# card 1
A:SR12-1>config>card# card-type iom3-xp
```

The 7450 ESS-7/12, and 7750 SR-7/12, and 7750 SR-12e platforms also support a variety of IMM in designated chassis slots. IMM has integrated MDAs. The provisioning requirements depends on the generation of IMM that you use. Refer to the IMM Installation Guide for more information.

The 7750 SR-a platforms support IOM-a cards in dedicated chassis slots. The 7750 SR-a4 supports one physical IOM-a in slot 3. This IOM-a is represented in the CLI as card 1. The 7750 SR-a8 supports two physical IOM-a cards, one in slot 3, the other in slot 6. These IOM-a cards are represented in the CLI as card 1 and card 2 respectively. The IOM-a does not have pluggable MDA slots. Each IOM-a can be configured to support up to four MDA-a or MDA-aXP modules. IOM-a cards are configured in the same manner as IOMs.

The 7750 SR-e platforms support the IOM-e modules in dedicated slots in the rear of each chassis. The 7750 SR-1e supports one physical IOM-e module. This IOM-e is represented in the CLI as card 1. The 7750 SR-2e supports two physical IOM-e cards. These IOM-e cards are represented in the CLI as card 1 and card 2 respectively. The 7750 SR-3e supports three physical IOM-e cards. These IOM-e cards are represented in the CLI as card 1, card 2, and card 3 respectively. The IOM-e does not have pluggable MDA slots. An IOM-e can be configured to support up to four MDA-e modules. IOM-e cards are configured in the same manner as IOMs.

The 7750 SR-c4/c12 platforms do not have slots for IOM or IMM cards. The system is modeled as having a fixed system-provisioned IOM in slot 1. The chassis has positions that accept MCMs or CMAs. MCMs accept MDAs. CMAs can be directly inserted into the 7750 SR-c4/c12 without the need for MCMs. CMAs are modeled as MDAs in SR OS.

2.1.2 MCMs

MCMs are only supported on the 7750 SR-c12 and SR-c4 systems.

An MCM must be configured before an MDA can be provisioned. If you provision an MDA type before an MCM is configured, it is assumed you are provisioning a CMA. CMAs do not require MCM preconfiguration. Up to six MCMs may be provisioned on a 7750 SR-c12. Even-numbered CMA positions are invalid for MCM installation (MCMs physically span two CMA positions; “mcm 1” spans CMA position 1 and 2). Up to two MCMs can be provisioned on the 7750 SR-c4.

Refer to the CMA Installation Guide and MDA Installation Guide for more information on the physical characteristics of each card.

2.1.3 MDA-a, MDA-aXP, MDA, MDA-XP and MDA-e Modules

MDAs are pluggable adapter cards that provide physical interface connectivity. MDAs are available in a variety of interface and density configurations. MDA modules differ by chassis. Refer to the individual chassis guide and the individual MDA installation guides for more information about specific MDAs.

On the 7750 SR-c4/c12 platforms, the MDAs plug into MCMs. MCMs must be provisioned before an MDA can be provisioned with a type. Up to six MDAs (each seated in an MCM) may be provisioned on a 7750 SR-c12. Even-numbered CMA positions are invalid for MDA installation (MDAs physically span two CMA positions; “mda 1” spans CMA positions 1 and 2). Up to two MDAs (each seated in an MCM) can be provisioned on the 7750 SRc4. CMAs are also supported on 7750 SR-c4/c12 platforms, as described in [CMAs](#).

The following displays a **show card state** command. In this example, an **m60-10/100eth-tx** MDA is installed in position 1 on a 7750 SR-c12.

```
A:ALU-3>config>card# show card state
=====
Card State
=====
```

Slot/ Id	Provisioned Type	Equipped Type	Admin State	Operational State	Num Ports	Num MDA	Comments
1	iom-xp	iom-xp	up	up		12	
1/1	mcm-xp	mcm-xp	up	up			
1/3		mcm-xp	up	unprovisioned			
1/1	m60-10/100eth-tx	m60-10/100eth-tx	up	up			
1/5	c8-10/100eth-tx	c8-10/100eth-tx	up	up			
1/6		c1-1gb-sfp	up	unprovisioned			
1/7		c8-chds1	up	unprovisioned			
1/8		c4-ds3	up	unprovisioned			
1/9		c8-10/100eth-tx	up	unprovisioned			
1/10		c1-1gb-sfp	up	unprovisioned			
1/11		c8-chds1	up	unprovisioned			
1/12		c4-ds3	up	unprovisioned			
A	cfm-xp	cfm-xp	up	up			Active
B	cfm-xp		up	down			Standby

```
=====
A:ALU-3>config>card#
```

On the 7450 ESS-7/12, 7750 SR-7/12, and 7750 SR-12e, MDAs plug into IOMs. (MDA and MDA-XP modules plug into the IOM3-XP/-B/-C. MDA-e modules plug into the IOM4-e and IOM4-e-B). Up to two MDAs can be provisioned on an IOM.

IMMs are designed with fixed integrated media cards, which may require provisioning, depending on the generation of the IMM.

MDA-a and MDA-aXP modules are used in the 7750 SR-a and the MDA-e and ISA2 modules are used in the 7750 SR-e chassis. Up to four MDAs can be provisioned for each IOM.

In all cases, the card slot and IOM or IMM card-type must be provisioned before an MDA can be provisioned. A preprovisioned MDA slot can remain empty without interfering with services on populated equipment. When an MDA is installed and enabled, the system verifies that the MDA type matches the provisioned type. If the parameters do not match, the MDA remains offline.

On the 7750 SR-c12/SR-c4, 7450 ESS-7/12, 7750 SR-7/12, and 7750 SR-12e platforms, MDA names in the CLI start with the letter 'm' (for example, m10-1gb-xp-sfp).

The following example displays the **card**, **card-type**, **mda**, and **mda-type** command usage in the 7750 SR-7:

```
A:SR7>config# card 1
A:SR7>config>card# card-type iom3-xp
A:SR7>config>card# mda 1
A:SR7>config>card>mda# mda-type m60-10/100eth-tx
A:SR7>config>card>mda# exit
A:SR7>config>card# mda 2
A:SR7>config>card>mda# mda-type m10-1gb-sfp
A:SR7>config>card>mda# exit
```

The following example displays the configuration:

```
A:SR7# admin display-config
. . .
-----
echo "Card Configuration "
#-----
card 1
    card-type iom3-xp
    mda 1
        mda-type m60-10/100eth-tx
    exit
    mda 2
        mda-type m10-1gb-sfp
    exit
exit
-----
```

The 7750 SR-a4 and 7750 SR-a8 support only MDA-a and MDA-aXP modules, which are identified in the CLI with an “ma” prefix (for example, ma4-10gb-sfp+), or “max” prefix (for example, maxp10-10gb-sfp+). Likewise, the 7750 SR-1e, 7750 SR-2e, and 7750 SR-3e support only MDA-e modules, which are identified in the CLI with an “me” prefix, such as me1-100gb-cfp2.

The following example shows the **card**, **card-type**, **mda**, and **mda-type** command usage in the 7750 SR-1e:

```
A:SR1e>config# card 1
A:SR1e>config>card# card-type iom-e
A:SR1e>config>card# mda 1
A:SR1e>config>card>mda# mda-type me10-10gb-sfp+
A:SR1e>config>card>mda# exit
A:SR1e>config>card# mda 4
A:SR1e>config>card>mda# mda-type me1-100gb-cfp2
A:SR1e>config>card>mda# exit
```

The following example displays the configuration:

```

A:SR1e# admin display-config
. . .
-----
echo "Card Configuration"
#-----
    card 1
      card-type iom-e
      mda 1
        mda-type me10-10gb-sfp+
      exit
      mda 4
        mda-type me1-100gb-cfp2
      exit
    exit
  -----
A:SR1e#

```

2.1.4 XMAs/C-XMAs



Note: For consistency across platforms, XMAs are modeled in the system as MDAs, and unless specified otherwise, the term MDA is used generically in this document to refer to both MDAs and C-XMA/XMAs. When the term XMA is used, it refers to both XMAs and C-XMAs unless specified otherwise.

XMAs are supported on the 7950 XRS platforms. XMAs plug into XCMs. XCMs must be provisioned before an XMA can be provisioned with a type.

The XMA information must be configured before ports can be configured. After you configure the XCM, use the following CLI commands to provision XMAs.

A maximum of two XMAs can be configured on an XCM. The following example displays the card slot, card type, MDA slot, and MDA type command usage:

```

A:XRS20>config# card 1
A:XRS20>config>card# card-type xcm-x20
A:XRS20>config>card# mda 1
A:XRS20>config>card>mda# mda-type cx2-100g-cfp
A:XRS20>config>card>mda# power-priority-level 130
A:XRS20>config>card>mda# exit
A:XRS20>config>card# mda 2
A:XRS20>config>card>mda# mda-type cx20-10g-sfp
A:XRS20>config>card>mda# power-priority-level 135
A:XRS20>config>card>mda# exit

```

The following example displays the configuration:

```

A:XRS20# admin display-config
. . .
-----

```

```

echo "Card Configuration "
#-----
    card 1
        card-type xcm-x20
        mda 1
            mda-type cx2-100g-cfp
            power-priority-level 130
        exit
        mda 2
            mda-type cx20-10g-sfp
            power-priority-level 135
        exit
    exit
exit
-----
A:XRS20#

```

On the 7950 XRS, the **show card state** output displays an “x” in the name of the XMA and “cx” in the name of a C-XMA:

```

A:Dut-A# show card state
=====
Card State
=====
Slot/   Provisioned Type      Admin Operational   Num   Num  Comments
Id      Equipped Type (if different) State State           Ports MDA
-----
1       xcm-x20                    up    up                2
1/1     cx20-10g-sfp              up    up                20
1/2     cx20-10g-sfp              up    up                20
2       xcm-x20                    up    up                2
2/1     cx20-10g-sfp              up    up                20
A       cpm-x20                    up    up                Active
B       cpm-x20                    up    up                Standby
=====

```

2.1.5 Hardware Licensing

With the introduction of pay-as-you-grow licensing, FP4-based assemblies (IOMs and XMA) now include variants with license levels. These levels define the capacity and functionality of the assembly. The capacity controls aspects such as the number and types of connectors that can be configured, as well as the total connector bandwidth. Licensing also controls the number of hardware egress queues and egress policers that are available for services. For a more complete description of the levels available for a particular assembly, see the associated Installation Guide.

The license level must be provisioned for the assembly at the same time as the card type or MDA type is provisioned. Each assembly has a set of levels applicable to that particular IOM or XMA that are defined using mnemonic strings. For example, an assembly may have a level 'cr1200g' which refers to a functional level of 'core routing' and a capacity maximum bandwidth of 1.2 Tb/s. A second example of a license level is 'he2400g+', which refers to a functional level of 'high scale edge routing' and a capacity level of a bandwidth of 2.4 Tb/s but with Intelligent Fan In/Out to a higher bandwidth.

When an assembly is installed in the chassis, the license level encoded into the equipped assembly must match the value provisioned for the assembly. If they do not match, the assembly can not become active in the chassis. The only exception is that a variant of the assembly with the maximum functional and capacity level is allowed to come up in a slot provisioned as any level; the restrictions in effect will be at the provisioned level, but this allows this specific assembly to be used to replace any other level of that assembly if necessary.

The following example shows the provisioning of an XCM with two XMAs. The first XMA is a two complex, 2.4T 24-connector QSFP28 XMA with a license level of er2400g (edge routing, 2.4 Tb/s) and the second XMA is a two complex, 2.4T 6-connector CFP8 XMA with a license level of he1600g (high scale edge routing, 4 connector 1.2Tbps):

```
*A:bkvm20# configure card 4 card-type "xcm2-x20"
*A:bkvm20# configure card 4 mda 1 mda-type "x24-100g-qsfp28" level "er2400g"
*A:bkvm20# configure card 4 mda 2 mda-type "x6-400g-cfp8" level "he1600g"
*A:bkvm20# show mda
=====
MDA Summary
=====
```

Slot	Mda	Provisioned Type Equipped Type (if different)	Admin State	Operational State
4	1	x24-100g-qsfp28:er2400g	up	up
	2	x6-400g-cfp8:he1600g	up	up

```
=====
```

The **show card** and **show mda** output display both the type and level of the assembly and will indicate when there is a difference between the provisioned and installed levels. In the following example, the first XMA has a provisioned value matching the installed assembly and the second XMA has a difference in the provisioned and installed assembly.

```
*A:bkvm20# show mda 4/1 detail
=====
MDA 4/1 detail
=====
```

Slot	Mda	Provisioned Type Equipped Type (if different)	Admin State	Operational State
4	1	x24-100g-qsfp28:er2400g	up	up

```
MDA Licensing Data
  Licensed Level          : er2400g
  Description             : 2.4T, 24c, Edge Routing
*A:bkvm20# show mda 4/2 detail
=====
MDA 4/2 detail
=====
Slot  Mda  Provisioned Type                               Admin  Operational
        Equipped Type (if different)                     State   State
-----
  4     2     x6-400g-cfp8:he2400g                        up      provisioned
        x6-400g-cfp8:he1600g
MDA Licensing Data
  Licensed Level          : he1600g
  Description             : 1.6T, 4c, High Scale Edge Routing
```

2.1.6 Software License Activation

The pay-as-you-grow licensing of the 7750 SR hardware platforms includes the ability to distribute software based licenses to operational systems in a live network. This is done by the creation of a license key file containing various individual licenses, making this file available to a system, and then activating that license key file on the system. Nokia provides the Configuration License Manager (CLM) tool to assist in this process. Normally, the CLM would perform all the steps in license distribution to a target system and only the assignment of individual licenses within a system itself need to be managed using the management mechanism for the system itself. Refer to the *7x50 Hardware Platform License Guide* for more detail on the use of CLM.

The license key file is a secure distribution mechanism for the licenses. Within the license key file, are one or more license keys. Each license key is assigned for a particular target system identified using a UUID. This UUID is tied to the chassis of the system and therefore the license key can only be validated on the system with that chassis. Each license key is also tied to a specific major software release of SR OS.

When a license key file is available to a target system and is activated it is first validated to ensure it is applicable for the specified target. For the license key of the SR OS hardware platform to be valid, the operator must ensure that the:

- license is for a 7xxx platform
- UUID of the system matches the one encoded in the "UUID-locked" license key
- SR OS software version (the major release number) matches the one encoded in the license key
- license file is not expired

Validation or activation of the license key file will result in zero or one license key that is valid for the running software on the target. If there is a valid license key, the records contained in that key are read and made available to the system (see [Software License Records](#)).

There may be additional license keys in the file that are for the target system but for a different software release. This will be the case when an upgrade of the system is planned. Those license keys shall be considered available however, the feature licenses contained within them are not available for use. These can be seen using the **show system license available-licenses** command.

```
A:bkvm20# show system license available-licenses
=====

Available Licenses
License name   : sr-regress@list.nokia.com
License uuid   : ab516e50-2413-44aa-9f7c-34b4e5b64d19
Machine uuid   : ab516e50-2413-44aa-9f7c-34b4e5b64d19
License desc   : 7xxx Platform
License prod   : 7xxx Platform
License sros   : TiMOS-[BC]-16.0.*
Current date   : FRI NOV 03 15:53:54 UTC 2017
Issue date    : FRI SEP 22 20:55:14 UTC 2017
Start date     : FRI SEP 15 00:00:00 UTC 2017
End date      : THU MAR 15 00:00:00 UTC 2018
-----
License name   : sr-regress@list.nokia.com
License uuid   : ab516e50-2413-44aa-9f7c-34b4e5b64d19
Machine uuid   : ab516e50-2413-44aa-9f7c-34b4e5b64d19
License desc   : 7xxx Platform
License prod   : 7xxx Platform
License sros   : TiMOS-[BC]-17.0.*
Current date   : FRI NOV 03 15:53:54 UTC 2017
Issue date    : FRI SEP 22 20:55:14 UTC 2017
Start date     : FRI SEP 15 00:00:00 UTC 2017
End date      : THU MAR 15 00:00:00 UTC 2018
-----
2 license(s) available.
=====
```

On a system boot, the license key file pointed to by the BOF is activated. If new license key files are activated on a system, the BOF should be updated to point to the new license key file. If there are active licenses in system that are in use and the node reboots and the license key file pointed to by the BOF either fails to validate or does not contain the in-use license records, the system will be considered to be in unlicensed state. In this state, the node will reboot every 60 minutes until the records are no longer in-use or a valid license key is provided that includes all the in use license records.

2.1.7 Software License Records

The activate license key shall unlock a set of license records for use within the system. In release 16.0, only Hardware Upgrade license records are distributed in the license keys. These can be used to upgrade the hardware capacity or the hardware functional level of a card or an XMA. These upgrades define a starting level and an upgraded level for the target assembly. Multiple instances of the same upgrade can be assigned to a system in the license key. The list of these records and the number in-use and available can be checked with the **show licensing-entitlements** command.

```
*A:bkvm20# show licensing entitlements
```

```
=====
```

License	Available	In-Use	State

MDA Upgrades			
cr1200g-cr1600g	1	1	VALID
cr1200g-er1200g	1	0	VALID
cr1600g-cr2400g	1	0	VALID
cr1600g-er1600g	2	0	VALID
cr2400g-er2400g	1	0	VALID
er1200g-er1600g	4	2	VALID
er1200g-he1200g	1	0	VALID
er1600g-er2400g	1	0	VALID
er1600g-he1600g	1	0	VALID
er2400g-he2400g	1	0	VALID
he1200g-he1600g	1	0	VALID
he1600g-he2400g	1	0	VALID

```
=====
```

*

An upgrade can be assigned to a particular card or XMA using the **upgrade** command within the **card** or **mda** context. Multiple upgrades can be assigned to the same card or XMA if desired to gradually augment the base card. In the following example, four upgrades have been applied in sequence to the base level of cr1600g to bring the XMA from CR to ER then ER to HE then to increase capacity first from 1600 Gbps to 2400Gbps and then from 2400 Gbps to 2400 Gbps with aggregation to 3600 Gbps.

```
=====
```

MDA 1/1 detail

```
=====
```

Slot	Mda	Provisioned Type Equipped Type (if different)	Admin State	Operational State

1	1	s36-100gb-qsfp28:cr1600g (not equipped)	up	provisioned

```
=====
```

MDA Licensing Data

Licensed Level	: he2400g+
Description	: 2.4T /w agg to 3.6T, 36p, High Scale Edge

	Routing
Upgrade 1	: cr1600g-er1600g
Upgrade 2	: er1600g-he1600g
Upgrade 3	: he1600g-he2400g
Upgrade 4	: any2400g-2400g+

2.1.8 CMAs

CMAs (Compact Media Adapter) are supported on the 7750 SR-c12 and SR-c4 and are configured and provisioned in the same manner as MDAs. Up to twelve CMAs may be provisioned on a 7750 SR-c12, and up to four CMAs may be provisioned on an SR-c4. CMA names in the CLI start with the letter “c” (for example, c4-ds3). The following shows **show card state** command output. In this example, a **c8-10/100eth-tx** CMA is installed in CMA position 5.

```
A:7750-3# show card state
=====
Card State
=====
```

Slot/ Id	Provisioned Type	Equipped Type	Admin State	Operational State	Num Ports	Num MDA	Comments
1	iom-xp	iom-xp	up	up		12	
1/5	c8-10/100eth-tx	c8-10/100eth-tx	up	up	8		
1/6	c8-10/100eth-tx	c8-10/100eth-tx	up	up	8		
1/7	c8-chds1		up	unprovisioned			
1/8	c4-ds3		up	unprovisioned			
1/9	c8-10/100eth-tx		up	unprovisioned			
1/10	c1-lgb-sfp		up	unprovisioned			
1/11	c8-chds1		up	unprovisioned			
1/12	c4-ds3		up	unprovisioned			
A	cfm-xp	cfm-xp	up	up			Active
B	cfm-xp		up	provisioned			Standby

```
=====
A:7750-3#
```

On the 7750 SR-c4 platform there are two fixed 10GbE ports that are modeled in CLI as an icm2-10gb-xp-xfp (integrated CMA) in position 1/5. A preprovisioned CMA slot can remain empty without conflicting with populated slots.

Once installed and enabled, the system verifies that the installed CMA type matches the provisioned type. If the parameters do not match, the CMA remains offline.

2.1.9 Versatile Service Module (VSM)

The Versatile Service Module (VSM) is a module that allows operators to internally connect a VPLS or VLL service into an IES or IPVPN service. Each module is capable of 10 Gb/s throughput.

A VSM, like an MDA, is installed and provisioned as a pluggable module in an IOM.

The VSM is supported on the 7450 ESS-7/12, 7750 SR-7/12, and 7750 SR-12e platforms. The VSM is not supported on the 7950 XRS or on the 7750 SR-c12/c4 platforms.

See [Versatile Service Module](#) for more details.

2.1.10 Oversubscribed Ethernet MDAs

The 7750 SR and 7450 ESS support oversubscribed Ethernet MDAs and CMAs. These have more bandwidth towards the user than the capacity between the MDA and IOM.

A traffic management function is implemented on the MDA to control the data entering the IOM. This function consists of two parts:

- rate limiting
- packet classification and scheduling

2.1.10.1 Rate Limiting

The oversubscribed MDA or CMA limits the rate at which traffic can enter the MDA or CMA on a per-port basis. If a port exceeds its configured limits then the excess traffic will be discarded, and 802.3x flow control frames (pause frames) are generated.

2.1.10.2 Packet Classification and Scheduling

The classification and scheduling function implemented on the oversubscribed MDA or CMA ensures that traffic is correctly prioritized when the bus from the MDA or CMA to the IOM is overcommitted. This could occur if the policing parameters configured are such that the sum of the traffic being admitted into the MDA or CMA is greater than the capacity between the MDA and the IOM.

The classification function uses the bits set in the DSCP or Dot1p fields of the customer packets to perform classification. It can also identify locally addressed traffic arriving on network ports as Network Control packets. This classification on the oversubscribed MDA or CMA uses the following rules:

- If the service QoS policy for the SAP (port or VLAN) uses the default classification policy, all traffic is classified as Best Effort (be).
- If the service QoS policy for the SAP contains a Dot1p classification, the Dot1p field in the customer packets is used for classification on the MDA or CMA.
- If the service QoS policy for the SAP contains a DSCP classification, the DSCP field in the customer packets is used for classification on the MDA or CMA.
- If a mix of Dot1p and DSCP classification definitions are present in the service QoS policy, then the field used to perform classification is the type used for the highest priority definition. For example, if High Priority 1 is the highest priority definition and it specifies that the DSCP field should be used, then the DSCP field is used for classification on the MDA or CMA and the Dot1p field is ignored.
- If the service QoS policy for the SAP specifies IP or MAC filters for forwarding class identification, then traffic is treated as Best Effort. Full MAC or IP classification is not possible on the MDA/CMA (but is possible on the IOM).
- The packet is classified into 16 classes. Typically, these are the eight forwarding classes and each packet is assigned one priority per forwarding class. After classification, the packet is offered to the queuing model. This queuing model is limited to three queues each having four thresholds. These thresholds define whether an incoming packet, after classification, is accepted in the queue or not.

[Table 3](#) shows typical mapping of classes onto queues and thresholds.

Table 3 Typical Mapping Of Classes Onto Queues/Threshold

Counter	{Queue	Threshold	Traffic Class}
0	{2	3	"fc-nc / in-profile"
1	{2	2	"fc-nc / out-profile"
2	{2	1	"fc-h1 / in-profile"
3	{2	0	"fc-h1 / out-profile"

Table 3 Typical Mapping Of Classes Onto Queues/Threshold

Counter	{Queue	Threshold	Traffic Class}
4	{1	3	"fc-ef / in-profile"}
5	{1	2	"fc-ef / out-profile"}
6	{1	1	"fc-h2 / in-profile"}
7	{1	0	"fc-h2 / out-profile"}
8	{0	3	"fc-l1 / in-profile"}
9	{0	3	"fc-l1 / out-profile"}
10	{0	2	"fc-af / in-profile"}
11	{0	2	"fc-af / out-profile"}
12	{0	1	"fc-l2 / in-profile"}
13	{0	1	"fc-l2 / out-profile"}
14	{0	0	"fc-be / in-profile"}
15	{0	0	"fc-be / out-profile"}

A counter is associated with each mapping. The above is an example and is dependent on the type of classification (such as dscp-exp, dot1p, and so on). When the threshold of a particular class is reached, packets belonging to that class are not accepted in the queue. The packets are dropped and the associated counter is incremented.

The scheduling of the three queues is done in a strict priority, highest priority basis is associated with queue 2. This means that scheduling is done at queue level, not on the class that resulted from the classification. As soon as a packet has been accepted by the queue there is no way to differentiate it from other packets in the same queue (for example, another classification result not exceeding its threshold). All packets queued in the same queue, have the same priority from a scheduling point of view.

2.1.11 Channelized MDA/CMA Support

2.1.11.1 Channelized DS-1/E-1 CMA

Each 8-port channelized DS-1/E-1 CMA supports channelization down to DS-0. Each 8-port channelized DS-1/E-1 CMA supports 64 channel groups. This MDA is supported on the 7750 SR-7/12 and 7750 SR-c4/c12 platforms. This CMA is supported on the 7750 SR-c4/c12 platforms.

2.1.11.2 Channelized DS-3/E-3 CMA

On the E-3 CMA, bit stuffing is not supported in G.751 framing mode. All of the 12 justification service bits and the four justification bits contain valid data on the transmitted signal. Incoming bitstreams should contain valid data in the 12 justification service bits and four justification bits, otherwise the link will not function.

This CMA is supported on the 7750 SR-c4/c12 platforms.

2.1.11.3 Channelized Any Service Any Port (ASAP) CHOC-3/STM-1

Each port for the channelized ASAP OC-3/STM-1 MDA supports channelization down to DS-0 and accepts one OC-3/STM-1 SFP small form factor pluggable (SFP) module. The same SFP optics used on Nokia's SONET/SDH MDAs can be used on the channelized ASAP OC-3/STM-1 MDA.

Each channelized OC-3/STM-1 supports up to 512 channels with DS-0 timeslots with per-channel encapsulation configuration (for example, Frame Relay, PPP, cHDL, ATM). DS-3 TDM channels can be further channelized to DS-1/E-1 channel groups. An E3 TDM channel cannot be channelized and can only be configured in clear channel operation. The MDA is based on a programmable data path architecture that enables enhanced Layer 1 and Layer 2 data path functionality, for example ATM TM features, MDA-based channel or port queuing, or multilink applications like Inverse ATM Multiplexing (IMA). This MDA is supported on the 7750 SR-7/12 and the 7750 SR-c4/c12 platforms.

2.1.11.4 Channelized OC-12/STM-4 ASAP MDAs

The channelized OC-12/STM-4 variant of the ASAP MDAs has features and channelization options similar to the 4-port channelized OC-3/STM-1 ASAP MDA.

DS-3 TDM channels can be further channelized to DS-1/E-1 channel groups. An E-3 TDM channel cannot be channelized and can only be configured in clear channel operation. This MDA is supported on the 7750 SR-7/12 and the 7750 SR-c4/c12 platforms.

2.1.11.5 Channelized DS-3/E-3 ASAP MDA (4-Port)

The 4-port MDA provides four ports configurable as DS-3 or E-3. The MDA has eight (8) 1.0/2.3 connectors and accepts up to eight (8) DS-3/E-3 coax patch cables.

Each physical DS-3 connection can support a full clear-channel DS-3, or it can be channelized into independent DS-1/E-1 data channels. Each DS-1/E-1 channel can then be further channelized down to DS-0s. All DS-0 channels within a DS-3 port must be configured for the same channel speed.: 56 kb/s or 64 kb/s. The 56 kb/s speed value is only supported on DS-1 channels (ESF and SF framing) and not on E-1 (G.704) channels. Also, 56 kb/s channels cannot be part of a bundle. E-3 ports do not support channelization, only clear channel operation. This MDA is supported on the 7750 SR-7/12 and the 7750 SR-c4/c12 platforms. This MDA is supported on the 7750 SR-7/12 and the 7750 SR-c4/c12 platforms.

2.1.11.6 Channelized DS-3/E-3 ASAP MDA (12-Port)

The 12-port MDA provides 12 ports configurable as DS-3 or E-3. The MDA has 24 1.0/2.3 connectors and accepts up to 24 DS-3/E-3 coax patch cables.

Each physical DS-3 connection can support a full clear-channel DS-3, or it can be channelized into independent DS-1/E-1 data channels. Each DS-1/E-1 channel can then be further channelized down to DS-0s. All DS-0 channels within a DS-3 port must be configured for the same channel speed.: 56 kb/s or 64 kb/s. The 56 kb/s speed value is only supported on DS-1 channels (ESF and SF framing) and not on E-1 (G.704) channels. Also, 56 kb/s channels cannot be part of a bundle. E-3 ports do not support channelization, only clear channel operation. This MDA is supported on the 7750 SR-7/12 and the 7750 SR-c4/c12 platforms.

2.1.11.7 Channelized OC-3/STM-1 Circuit Emulation Services (CES) CMA and MDA

The channelized OC-3/STM-1/OC-12/STM-4 CES MDAs (c1-choc3-ces-sfp / m1-choc3-ces-sfp, m4-choc3-ces-sfp, m1-choc12-ces-sfp) provide an industry leading consolidation for DS-1, E-1 and n*64 kb/s for CES.

The channelized OC-3/STM-1/OC-12/STM-4 CES CMA/MDAs support CES. Circuit emulation services are interoperable with the existing 7705 SAR and 7250 SAS circuit emulation services. They are also interoperable with the 1850 TSS-5 circuit emulation services.

Two modes of circuit emulation are supported: unstructured and structured. Unstructured mode is supported for DS-1 and E-1 channels as per RFC4553 (SAToP). Structured mode is supported for n*64 kb/s circuits as per RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*. In addition, DS-1, E-1 and n*64 kb/s circuits are also supported as per MEF8, *Circuit Emulation Services over Ethernet (CESoETH)* (Oct 2004). TDM circuits are optionally encapsulated in MPLS or Ethernet as per the applicable standards.

All channels on the CES CMA/MDA are supported as circuits to be emulated across the packet network. This includes DS-1, E-1 and n*64 kb/s channels. Structure agnostic mode is supported for DS-1 and E-1 channels. Structure aware mode is supported for n*64 kb/s channel groups in DS-1 and E-1 carriers. N*64 kb/s circuit emulation supports basic and Channel Associated Signaling (CAS) options. CAS configuration must be identical for all channel groups on a given DS-1 or E-1.

Circuits encapsulated in MPLS, use circuit pipes (Cpipes) to connect to the far end circuit. Cpipes support either SAP-spoke SDP or SAP-SAP connections.

Circuits encapsulated in Ethernet can be selected as a SAP in Epipe. Circuits encapsulated in Ethernet can be either SAP-spoke SDP or SAP-SAP connections for all valid Epipe SAPs. An EC-ID and far-end destination MAC address must be configured for each circuit.

Each OC-3/STM-1 port can be independently configured to be loop-timed or node-timed. Each OC-3/STM-1 port can be configured to be a timing source for the node. Each DS-1 or E-1 channel can be independently configured to be loop-timed, node-timed, adaptive-timed, or differential-timed. One adaptive timed circuit is supported per CMA or MDA. The CES circuit configured for adaptive timing can be configured to be a timing source for the node. This is required to distribute network timing to network elements which only have packet connectivity to network.

On the 7750 SR-c12 CES CMA, a BITS port is also provided. The BITS port can be configured as one reference sources (ref1, ref2) in the system timing subsystem. These MDAs are supported on the 7750 SR-7/12 and the 7750 SR-c4/c12 platforms.

2.1.11.8 Network Interconnections

Nokia routers can fill the needs of smaller service providers as well as the more remote point of presence (PoPs) locations for larger service providers. To support the use of lower speed links as network links in the likelihood that lower speed circuits are used as network or backbone links, the routers support a DS-1/E-1/DS-3/E-3 port (ASAP MDAs) or channel and an MLPPP bundle (ASAP MDAs) as network ports to transport and forwarding of all service types. This feature allows service providers to use lower speed circuits to interconnect small PoPs and CoS that do not require large amounts of network or backbone bandwidth.

2.2 Digital Diagnostics Monitoring

Some Nokia SFPs, XFPs, QSFPs, CFPs and the MSA DWDM transponder have the Digital Diagnostics Monitoring (DDM) capability where the transceiver module maintains information about its working status in device registers including:

- temperature
- supply voltage
- transmit (TX) bias current
- TX output power
- received (RX) optical power

For QSFPs and CFPs, DDM Temperature and Supply voltage is available only at the Module level as shown in [Table 5](#).

Refer to the [Statistics Collection](#) section for details about the QSFP and CFP sample DDM and DDM Lane information.

For the QSFPs and CFPs, the number of lanes is indicated by DDM attribute “Number of Lanes: 4”.

Subsequently, each lane threshold and measured values are shown per lane.

If a given lane entry is not supported by the given QSFP or CFP specific model, then it is shown as “-” in the entry.

A sample QSFP and CFP lane information is provided below:

```
Transceiver Data
Transceiver Type   : QSFP+
Model Number      : 3HE06485AAAA01  ALU  IPUIBMY3AA
TX Laser Wavelength: 1310 nm
Number of Lanes    : 4
Connector Code     : LC
Vendor OUI         : e4:25:e9
Manufacture date   : 2012/02/02
Media              : Ethernet
Serial Number      : 12050188
Part Number        : DF40GELR411102A
Optical Compliance : 40GBASE-LR4
Link Length support: 10km for SMF

=====
Transceiver Digital Diagnostic Monitoring (DDM)
=====
                        Value High Alarm  High Warn   Low Warn   Low Alarm
-----
Temperature (C)        +35.6          +75.0        +70.0        +0.0        -5.0
Supply Voltage (V)      3.23           3.60         3.50         3.10         3.00
=====
Transceiver Lane Digital Diagnostic Monitoring (DDM)
```



```
=====
```

		High Alarm	High Warn	Low Warn	Low Alarm
Lane Tx Bias Current (mA)		78.0	75.0	25.0	20.0
Lane Rx Optical Pwr (avg dBm)		2.30	2.00	-11.02	-13.01

```
-----
```

Lane ID	Temp(C)/Alm	Tx Bias(mA)/Alm	Tx Pwr(dBm)/Alm	Rx Pwr(dBm)/Alm
1	-	43.5	-	0.42
2	-	46.7	-	-0.38
3	-	37.3	-	0.55
4	-	42.0	-	-0.52

```
=====
```

```
Transceiver Type      : CFP
Model Number         : 3HE04821ABAA01  ALU  IPUIBHJDAA
TX Laser Wavelength: 1294 nm
Number of Lanes      : 4
Connector Code       : LC
Vendor OUI           : 00:90:65
Manufacture date     : 2011/02/11
Media                 : Ethernet
Serial Number        : C22CQYR
Part Number          : FTLC1181RDNL-A5
Optical Compliance   : 100GBASE-LR4
Link Length support: 10km for SMF
```

Transceiver Digital Diagnostic Monitoring (DDM)

```
=====
```

	Value	High Alarm	High Warn	Low Warn	Low Alarm
Temperature (C)	+48.2	+70.0	+68.0	+2.0	+0.0
Supply Voltage (V)	3.24	3.46	3.43	3.17	3.13

```
=====
```

Transceiver Lane Digital Diagnostic Monitoring (DDM)

```
=====
```

		High Alarm	High Warn	Low Warn	Low Alarm
Lane Temperature (C)		+55.0	+53.0	+27.0	+25.0
Lane Tx Bias Current (mA)		120.0	115.0	35.0	30.0
Lane Tx Output Power (dBm)		4.50	4.00	-3.80	-4.30
Lane Rx Optical Pwr (avg dBm)		4.50	4.00	-13.00	-16.00

```
-----
```

Lane ID	Temp(C)/Alm	Tx Bias(mA)/Alm	Tx Pwr(dBm)/Alm	Rx Pwr(dBm)/Alm
1	+47.6	59.2	0.30	-10.67
2	+43.1	64.2	0.27	-10.31
3	+47.7	56.2	0.38	-10.58
4	+51.1	60.1	0.46	-10.37

```
=====
```

The transceiver is programmed with warning and alarm thresholds for low and high conditions that can generate system events. These thresholds are programmed by the transceiver manufacturer.

There are no CLI commands required for DDM operations, however, the **show>port port-id detail** command displays DDM information in the Transceiver Digital Diagnostics Monitoring output section.

DDM information is populated into the router's MIBs, so the DDM data can be retrieved by Network Management using SNMP. Also, RMON threshold monitoring can be configured for the DDM MIB variables to set custom event thresholds if the factory-programmed thresholds are not at the desired levels.

The following are potential uses of the DDM data:

- Optics degradation monitoring — With the information returned by the DDM-capable optics module, degradation in optical performance can be monitored and trigger events based on custom or the factory-programmed warning and alarm thresholds.
- Link or router fault isolation — With the information returned by the DDM-capable optics module, any optical problem affecting a port can be quickly identified or eliminated as the potential problem source.

Supported real-time DDM features are summarized in [Table 4](#).

Table 4 Real-Time DDM Information

Parameter	User Units	SFP/XFP Units	SFP	XFP	MSA DWDM
Temperature	Celsius	C	Supported	Supported	Supported
Supply Voltage	Volts	μ V	Supported	Supported	Not supported
TX Bias Current	mA	μ A	Supported	Supported	Supported
TX Output Power	dBm (converted from mW)	mW	Supported	Supported	Supported
RX Received Optical Power ⁴	dBm (converted from dBm) (Avg Rx Power or OMA)	mW	Supported	Supported	Supported
AUX1	parameter dependent (embedded in transceiver)	-	Not supported	Supported	Not supported
AUX2	parameter dependent (embedded in transceiver)	-	Not supported	Supported	Not supported

The factory-programmed DDM alarms and warnings that are supported are summarized in [Table 5](#).

Table 5 DDM Alarms and Warnings

Parameter	SFP/XFP Units	SFP	XFP	Required?	MSA DWDM
Temperature - High Alarm - Low Alarm - High Warning - Low Warning	C	Yes	Yes	Yes	Yes
Supply Voltage - High Alarm - Low Alarm - High Warning - Low Warning	μ V	Yes	Yes	Yes	No
TX Bias Current - High Alarm - Low Alarm - High Warning - Low Warning	μ A	Yes	Yes	Yes	Yes
TX Output Power - High Alarm - Low Alarm - High Warning - Low Warning	mW	Yes	Yes	Yes	Yes
RX Optical Power - High Alarm - Low Alarm - High Warning - Low Warning	mW	Yes	Yes	Yes	Yes
AUX1 - High Alarm - Low Alarm - High Warning - Low Warning	parameter dependent (embedded in transceiver)	No	Yes	Yes	No
AUX2 - High Alarm - Low Alarm - High Warning - Low Warning	parameter dependent (embedded in transceiver)	No	Yes	Yes	No

2.2.1 SFPs and XFPs

The availability of the DDM real-time information and the warning and alarm status is based on the transceiver. It may or may not indicate that DDM is supported. Although some Nokia SFPs support DDM, Nokia has not required DDM support in releases prior to Release 6.0. Non-DDM and DDM-supported SFPs are distinguished by a specific ICS value.

For SFPs that do not indicate DDM support in the ICS value, DDM data is available although the accuracy of the information has not been validated or verified.

For non-Nokia transceivers, DDM information may be displayed, but Nokia is not responsible for formatting, accuracy, and so on.

2.2.2 Statistics Collection

The DDM information and warnings and alarms are collected at one minute intervals, so the minimum resolution for any DDM events when correlating with other system events is one minute.

In the Transceiver Digital Diagnostic Monitoring section of the **show port *port-id* detail** command output:

- if the present measured value is higher than either or both of the High Alarm and High Warn thresholds, an exclamation mark “!” displays along with the threshold value
- if the present measured value is lower than either or both of the Low Alarm and Low Warn thresholds, an exclamation mark “!” displays along with the threshold value

```
B:SR7-101# show port 2/1/6 detail
.....
=====
Transceiver Digital Diagnostic Monitoring (DDM), Internally Calibrated
=====
      Value High Alarm  High Warn   Low Warn  Low Alarm
-----
Temperature (C)      +33.0+98.0   +88.0      -43.0-45.0
Supply Voltage (V)      3.31 4.12    3.60      3.00 2.80
Tx Bias Current (mA) 5.7 60.0    50.00.1  0.0
Tx Output Power (dBm)  -5.45 0.00   -2.00     -10.50  -12.50
Rx Optical Power (avg dBm) -0.65-3.00! -4.00!    -19.51  -20.51
=====
```

2.3 Ports

2.3.1 Port Types

Before a port can be configured, the slot must be provisioned with a card type and MDA type.

Nokia routers support the following port types:

- Ethernet — Supported Ethernet port types include:
 - Fast Ethernet (10/100BASE-T)
 - Gb Ethernet (1GbE, 1000BASE-T)
 - 10 Gb Ethernet (10GbE, 10GBASE-X)
 - 40 Gb Ethernet (40GbE)
 - 100 Gb Ethernet (100GbE)

Router ports must be configured as either access, hybrid, or network. The default is network.

- Access ports — Configured for customer facing traffic on which services are configured. If a Service Access Port (SAP) is to be configured on the port or channel, it must be configured as an access port or channel. When a port is configured for access mode, the appropriate encapsulation type must be configured to distinguish the services on the port or channel. Once a port has been configured for access mode, one or more services can be configured on the port or channel depending on the encapsulation value.
- Network ports — Configured for network-facing traffic. These ports participate in the service provider transport or infrastructure network. Dot1q is supported on network ports.
- Hybrid ports — Configured for access and network-facing traffic. While the default mode of an Ethernet port remains network, the mode of a port cannot be changed between the access, network, and hybrid values unless the port is shut down and the configured SAPs or interfaces are deleted. Hybrid ports allow a single port to operate in both access and network modes. The MTU of a port in hybrid mode is the same as in network mode, except for the 10/100 MDA. The default encapsulation for hybrid port mode is dot1q; it also supports QinQ encapsulation on the port level. Null hybrid port mode is not supported. Once the port is changed to hybrid, the default MTU of the port is changed to match the value of 9212 bytes currently used in network mode (higher than an access port). This is to ensure that both SAP and network VLANs can be accommodated. The only exception is when the port is a 10/100 Fast Ethernet. In those cases, the MTU in hybrid mode is set to 1522 bytes, which corresponds

to the default access MTU with QinQ, which is larger than the network dot1q MTU or access dot1q MTU for this type of Ethernet port. The configuration of all parameters in access and network contexts continues to be done within the port using the same CLI hierarchy as in existing implementation. The difference is that a port configured in mode hybrid allows both ingress and egress contexts to be configured concurrently. An Ethernet port configured in hybrid mode can have two values of encapsulation type: dot1q and QinQ. The NULL value is not supported since a single SAP is allowed, and can be achieved by configuring the port in the access mode, or a single network IP interface is allowed, which can be achieved by configuring the port in network mode. Hybrid mode can be enabled on a LAG port when the port is part of a single chassis LAG configuration. When the port is part of a multi-chassis LAG configuration, it can only be configured to access mode since MC-LAG is not supported on a network port and consequently is not supported on a hybrid port. The same restriction applies to a port that is part of an MC-Ring configuration.

For a hybrid port, the amount of the allocated port buffers in each of ingress and egress is split equally between network and access contexts using the following **config>port>hybrid-buffer-allocation>ing-weight access access-weight** [0 to 100] **network network-weight** [0 to 100] and **config>port>hybrid-buffer-allocation>egr-weight access access-weight** [0 to 100] **network network-weight** [0 to 100] commands.

Adapting the terminology in buffer-pools, the port's access active bandwidth and network active bandwidth in each ingress and egress are derived as follows (egress formulas shown only):

- $\text{total-hybrid-port-egress-weights} = \text{access-weight} + \text{network-weight}$
- $\text{hybrid-port-access-egress-factor} = \text{access-weight} / \text{total-hybrid-port-egress-weights}$
- $\text{hybrid-port-network-egress-factor} = \text{network-weight} / \text{total-hybrid-port-egress-weights}$
- $\text{port-access-active-egress-bandwidth} = \text{port-active-egress-bandwidth} \times \text{hybrid-port-access-egress-factor}$
- $\text{port-network-active-egress-bandwidth} = \text{port-active-egress-bandwidth} \times \text{hybrid-port-network-egress-factor}$

When a named pool policy is applied to the hybrid port's MDA or to the hybrid port, the port's fair share of total buffers available to the MDA is split into three parts: default pools, named pools local to the port, and named pools on the ports MDA. This allocation can be altered by entering the corresponding values in the **port-allocation-weights** parameter.

- WAN PHY — 10 G Ethernet ports can be configured in WAN PHY mode (using the **ethernet xgig** config). When configuring the port to be in WAN mode, you can change certain SONET/SDH parameters to reflect the SONET/SDH requirements for this port.

- SONET-SDH and TDM — Supported SONET-SDH and TDM port types include:
 - n*DS-0 inside DS-1/E-1
 - DS-1/E-1DS-3/E-3
 - OC3/STM-1
 - OC12/STM-4
 - OC48/STM-16
 - OC192/STM-64 SONET/SDH
 - OC768/STM-256

A SONET/SDH port/path or a TDM port/channel can be configured with the following encapsulations depending on the MDA type:

- Frame Relay
 - PPP
 - cHDLC
- ATM — Some MDAs support ATM encapsulation on SONET/SDH and TDM ports. The ATM cell format can be configured for either UNI or NNI cell format. The format is configurable on a SONET/SDH or TDM port/channel path basis. All VCs on a path, channel or port must use the same cell format. The ATM cell mapping can also be configured on per-interface basis for either Direct or PLCP on some MDAs (for example ASAP MDA).
- Several Media Dependent Adapters (MDAs) support channelization down to the DS-0 level. ATM, Frame Relay, PPP, and cHDLC are supported encapsulations on channelized ports.
- Link Aggregation (LAG) — LAG can be used to group multiple ports into one logical link. The aggregation of multiple physical links allows for load sharing and offers seamless redundancy. If one of the links fails, traffic is redistributed over the remaining links.
- Multilink Bundles — A multilink bundle is a collection of channels on channelized ports that physically reside on the same MDA. Multilink bundles are used by providers who offer either bandwidth-on-demand services or fractional bandwidth services (fraction of a DS-3/E-3 for example). Multilink bundles are supported over PPP channels (MLPPP) and ATM channels (IMA).
- APS — Automatic Protection Switching (APS) is a means to provide redundancy on SONET equipment to guard against linear unidirectional or bidirectional failures. The network elements (NEs) in a SONET/SDH network constantly monitor the health of the network. When a failure is detected, the network proceeds through a coordinated pre-defined sequence of steps to transfer (or switchover) live traffic to the backup facility (called protection facility.) This is done very quickly to minimize lost traffic. Traffic remains on the protection facility until the primary facility (called working facility) fault is cleared, at which time the traffic may optionally be reverted to the working facility.

- **Bundle Protection Group (BPGRP)** — A BPGRP is a collection of two bundles created on the APS Group port. Working bundle resides on the working circuit of the APS group, while protection bundle resides on the protection circuit of the APS group. APS protocol running on the circuits of the APS Group port monitors the health of the SONET/SDH line and based on it or administrative action moves user traffic from one bundle to another in the group as part of an APS switch.
- **Cross connect adapter (CCA)** — A CCA on a VSM module interconnects the egress forwarding path on the IOM directly to the ingress forwarding path. This eliminates the need for the physical port MAC, PHY, cable and other MDA-specific components producing a less costly and more reliable adapter.
- **Optical Transport Network (OTN)** — Including OTU2, OTU2e, OTU3, and OTU4. OTU2 encapsulates 10-Gigabit Ethernet WAN and adds FEC (Forward Error Correction). OTU2e encapsulates 10-Gigabit Ethernet LAN and adds FEC (Forward Error Correction). OTU3 encapsulated OC768 and adds FEC. OTU4 encapsulates 100-Gigabit Ethernet and adds FEC.
- **Connector** — A QSFP28 (or QSFP-DD) connector that can accept transceiver modules including breakout connectors to multiple physical ports. For example, a QSFP28 connector can support four 10 Gb Ethernet ports. The connectors themselves cannot be used as ports in other commands, however, the breakout ports can be used as any Ethernet port.

2.3.2 Port Features

2.3.2.1 Port State and Operational State

There are two port attributes that are related and similar but have slightly different meanings: Port State and Operational State (or Operational Status).

The following descriptions are based on normal individual ports. Many of the same concepts apply to other objects that are modeled as ports in the router such as PPP/IMA/MLFR multilink bundles or APS groups but the show output descriptions for these objects should be consulted for the details.

- **Port State**
 - Displayed in port summaries such as **show port** or **show port 1/1**
 - tmnxPortState in the TIMETRA-PORT-MIB
 - Values: None, Ghost, Down (linkDown), Link Up, Up
- **Operational State**

- Displayed in the show output of a specific port such as **show port 2/1/3**
- tmnxPortOperStatus in the TIMETRA-PORT-MIB
- Values: Up (inService), Down (outOfService)

The behavior of Port State and Operational State are different for a port with link protocols configured (Eth OAM, Eth CFM or LACP for Ethernet ports, LCP for PPP/POS ports). A port with link protocols configured only transitions to the **Up** Port State when the physical link is up and all the configured protocols are up. A port with no link protocols configured transitions from Down to Link Up and then to Up immediately once the physical link layer is up.

The linkDown and linkUp log events (events 2004 and 2005 in the SNMP application group) are associated with transitions of the port Operational State. Note that these events map to the RFC 2863, *The Interfaces Group MIB*, (which obsoletes RFC 2233, *The Interfaces Group MIB using SMIv2*) linkDown and linkUp traps as mentioned in the SNMPv2-MIB.

An Operational State of **Up** indicates that the port is ready to transmit service traffic (the port is physically up and any configured link protocols are up). The relationship between port Operational State and Port State is shown in [Table 6](#):

Table 6 Relationship of Port State and Oper State

	Operational State (Oper State or Oper Status) (as displayed in “show port x/y/z”)	
Port State (as displayed in the show port summary)	For ports that have no link layer protocols configured	For ports that have link layer protocols configured (PPP, LACP, 802.3ah EFM, 802.1ag Eth-CFM)
Up	Up	Up
Link Up (indicates the physical link is ready)	Up	Down
Down	Down	Down

2.3.2.2 802.1x Network Access Control

Nokia routers support network access control of client devices (PCs, STBs, and so on) on an Ethernet network using the IEEE. 802.1x standard. 802.1x is known as Extensible Authentication Protocol (EAP) over a LAN network or EAPOL.

2.3.2.2.1 802.1x Modes

Nokia routers support port-based network access control for Ethernet ports only. Every Ethernet port can be configured to operate in one of three different operation modes, controlled by the port-control parameter:


- **force-auth** — Disables 802.1x authentication and causes the port to transition to the authorized state without requiring any authentication exchange. The port transmits and receives normal traffic without requiring 802.1x-based host authentication. This is the default setting.
- **force-unauth** — Causes the port to remain in the unauthorized state, ignoring all attempts by the hosts to authenticate. The switch cannot provide authentication services to the host through the interface.
- **auto** — Enables 802.1x authentication. The port starts in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. Both the router and the host can initiate an authentication procedure as described below. The port remains in unauthorized state (no traffic except EAPOL frames is allowed) until the first client is authenticated successfully. After this, traffic is allowed on the port for all connected hosts.

2.3.2.2.2 802.1x Basics

The IEEE 802.1x standard defines three participants in an authentication conversation (see [Figure 1](#) which shows an example with the 7450 ESS).

- The supplicant — This is the end-user device that requests access to the network.
- The authenticator — Controls access to the network. Both the supplicant and the authenticator are referred to as Port Authentication Entities (PAEs).
- The authentication server — Performs the actual processing of the user information.


Figure 1 802.1x Architecture



The authentication exchange is carried out between the supplicant and the authentication server, the authenticator acts only as a bridge. The communication between the supplicant and the authenticator is done through the Extended Authentication Protocol (EAP) over LANs (EAPOL). On the back end, the communication between the authenticator and the authentication server is done with the RADIUS protocol. The authenticator is thus a RADIUS client, and the authentication server a RADIUS server.

The messages involved in the authentication procedure are shown in [Figure 2](#). The router initiates the procedure when the Ethernet port becomes operationally up, by sending a special PDU called EAP-Request/ID to the client. The client can also initiate the exchange by sending an EAPOL-start PDU, if it doesn't receive the EAP-Request/ID frame during bootup. The client responds on the EAP-Request/ID with a EAP-Response/ID frame, containing its identity (typically username + password).

Figure 2 802.1x Authentication Scenario



OSSG039


After receiving the EAP-Response/ID frame, the router will encapsulate the identity information into a RADIUS AccessRequest packet, and send it off to the configured RADIUS server.

The RADIUS server checks the supplied credentials, and if approved will return an Access Accept message to the router. The router notifies the client with an EAP-Success PDU and puts the port in authorized state.

2.3.2.2.3 802.1x Timers

The 802.1x authentication procedure is controlled by a number of configurable timers and scalars. There are two separate sets, one for the EAPOL message exchange and one for the RADIUS message exchange. See [Figure 3](#) for an example of the timers on the 7750 SR.

Figure 3 802.1x EAPOL Timers (left) and RADIUS Timers (right)



OSSG040-7750

EAPOL timers:

- **transit-period** — Indicates how many seconds the Authenticator will listen for an EAP-Response/ID frame. If the timer expires, a new EAP-Request/ID frame will be sent and the timer restarted. The default value is 60. The range is 1 to 3600 seconds.
- **supplicant-timeout** — This timer is started at the beginning of a new authentication procedure (transmission of first EAP-Request/ID frame). If the timer expires before an EAP-Response/ID frame is received, the 802.1x authentication session is considered as having failed. The default value is 30. The range is 1 to 300.

- **quiet-period** — Indicates number of seconds between authentication sessions. It is started after logoff, after sending an EAP-Failure message or after expiry of the supplicant-timeout timer. The default value is 60. The range is 1 to 3600.

RADIUS timer and scaler:

- **max-auth-req** — Indicates the maximum number of times that the router will send an authentication request to the RADIUS server before the procedure is considered as having failed. The default value is value 2. The range is 1 to 10.
- **server-timeout** — Indicates how many seconds the authenticator will wait for a RADIUS response message. If the timer expires, the access request message is sent again, up to *max-auth-req* times. The default value is 60. The range is 1 to 3600 seconds.

The router can also be configured to periodically trigger the authentication procedure automatically. This is controlled by the `enable re-authentication` and `reauth-period` parameters. `Reauth-period` indicates the period in seconds (since the last time that the authorization state was confirmed) before a new authentication procedure is started. The range of `reauth-period` is 1 to 9000 seconds (the default is 3600 seconds, one hour). Note that the port stays in an authorized state during the re-authentication procedure.

2.3.2.2.4 802.1x Tunneling

Tunneling of untagged 802.1x frames received on a port is supported for both Epipe and VPLS service using either null or default SAPs (for example 1/1/1.*) when the port `dot1x port-control` is set to `force-auth`.

When tunneling is enabled on a port (using the command `configure port port-id ethernet dot1x tunneling`), untagged 802.1x frames are treated like user frames and are switched into Epipe or VPLS services which have a corresponding null SAP or default SAP on that port. In the case of a default SAP, it is possible that other non-default SAPs are also present on the port. Untagged 802.1x frames received on other service types, or on network ports, are dropped.

When tunneling is required, it is expected that it is enabled on all ports into which 802.1x frames are to be received. The configuration of `dot1x` must be configured consistently across all ports in LAG as this is not enforced by the system.

Note that 802.1x frames are treated like user frames, that is, tunneled, by default when received on a spoke or mesh SDP.

2.3.2.2.5 802.1x Configuration and Limitations

Configuration of 802.1x network access control on the router consists of two parts:

- Generic parameters, which are configured under **config>security>dot1x**
- Port-specific parameters, which are configured under **config>port>ethernet>dot1x**

801.x authentication:

- Provides access to the port for any device, even if only a single client has been authenticated.
- Can only be used to gain access to a pre-defined Service Access Point (SAP). It is not possible to dynamically select a service (such as a VPLS service) depending on the 802.1x authentication information.
- If 802.1x access control is enabled and a high rate of 802.1x frames are received on a port, that port will be blocked for a period of 5 minutes as a DoS protection mechanism.

2.3.2.3 MACsec


Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. MACsec provides point-to-point and point-to-multipoint security on Ethernet links between directly-connected nodes or nodes connected via a Layer 2 cloud. It is capable of identifying and preventing most security threats, including:

- denial of service
- intrusion
- man-in-the-middle
- masquerading
- passive wiretapping
- playback attacks

MACsec is standardized in IEEE 802.1AE, and is a Layer 2 encryption. MACsec encrypts anything from the 802.1AE header to the end of the payload including 802.1Q. MACsec leaves the DMAC and SMAC in clear text.

[Figure 4](#) shows the 802.1AE LAN-Mode structure.

Figure 4 802.1AE LAN-MODE



The forwarding on a MACsec packet is performed using the destination MAC address, which is in clear text.

2.3.2.3.1 MACsec 802.1AE Header (SecTAG)

The 802.1AE header includes a security TAG which includes:


- the association number within the channel
- the packet number to provide a unique initialization vector for encryption and authentication algorithms as well as protection against replay attack
- an optional LAN-Wide secure channel identifier

The security TAG (SecTAG) is identified by the MACsec EtherType and conveys the following:

- TAG Control Information (TCI)
- Association Number (AN)
- Short Length (SL)
- Packet Number (PN)
- Optionally-encoded Secure Channel Identifier (SCI)

Figure 5 shows the format of the SecTAG.

Figure 5 SecTAG Format



2.3.2.3.2 MACsec Encryption Mode

There are two main modes of encryption in MACsec:


- VLAN in clear text (WAN Mode)
- VLAN encrypted

802.1AE dictates that the 802.1Q VLAN needs to be encrypted. Some vendors give the option of configuring the MACsec on a port with VLAN in clear text.

SR OS supports both modes. On the 7750 SR, 7450 ESS, and 7950 XRS, 1/10 Gig cards support both mode of operation.

Figure 6 shows the VLAN encrypted and VLAN in clear.

Figure 6 802.1 AE LAN/WAN Modes and VLAN Encrypted/Clear



sw0119

MACsec Encryption per Traffic Flow Encapsulation Matching

In Release 16.0 and later, MACsec can be applied to a selected sub-set of the port traffic, based on the type and value of the packet encapsulation. The SR OS can be configured to match and encrypt the following traffic encapsulation types:

- All encap traffic arriving on port including untagged, single-tag, and double-tag. This is the default behavior of MACsec and the only option supported in releases prior to 16.0.
- Untagged only traffic.
- Single-tag or dot1q traffic. In this mode, MACsec can apply to a specific tag or wild card tag where all single-tag traffic is matched.

- Double-tag or q-in-q traffic. In this mode, MACsec can apply to a specific service tag, a specific service and customer tag, or a wildcard for any q-in-q traffic.

MKA PDUs are generated specifically for the traffic encapsulation type that is being matched.

2.3.2.3.3 MACsec Key Management Modes

There are four main, key management modes in MACsec. [Table 7](#) describes these management modes.

Table 7 MACsec Key Management Modes


Keying	Explanation	SR OS Support	Where Used
Static SAK	Manually configure each node with a static SAK, SAM, or CLI	NA	Switch to switch
Static CAK PRE SHARED KEY	Uses a dynamic MACsec Key Management (MKA) and uses a configured pre shared key to drive the CAK. The CAK encrypts the SAK between two peers and authenticates the peers	Supported	Switch to switch
Dynamic CAK EAP Authentication	Uses a dynamic MKA and uses a EAP MSK (Master System Key) to drive the CAK. The CAK encrypts the SAK between two peers and authenticates the peers	Not Supported	Switch to switch

Table 7 **MACsec Key Management Modes (Continued)**

Keying	Explanation	SR OS Support	Where Used
Dynamic CAK MSK distribution via RADIUS and EAP-TLS	MSKs are stored in the Radius server and distributed to the hosts via EAP-TLS. This is usually used in the access networks where there are a large number of hosts using MACsec and connecting to an access switch. MKA uses MSK to drive the CAK. The CAK encrypts the SAK between 2 peers and authenticates the peers	Not Supported	Host to switch

2.3.2.3.4 MACsec Terminology

Figure 7 illustrates some of the main concepts used in MACsec for the static-CAK scenario.

Figure 7 **MACsec Concepts for Static-CAK**

sw0120

Table 8 describes the MACsec terminology.

Table 8 MACsec Terms

MACsec Term	Description
CA: Connectivity Association	A security relationship, established and maintained by key agreement protocols (MKA), that comprises a fully-connected subset of the service access points in stations attached to a single LAN that are to be supported by MACsec.
MKA: MACsec Key Agreement Protocol	Control protocol between MACsec peers, which is used for peer aliveness and encryption key distribution. MACsec Key Agreement is responsible for discovering, authenticating, and authorizing the potential participants in a CA.
SecY: MAC Security Entity	Operates the MAC Security protocol within a system. Manages and identifies the SC and the corresponding active SA.
SC: Security Channel	SC provides a unidirectional point-to-point or point-to-multipoint communication. Each SC contains a succession of SAs and each SC has a different SAK.
SA: Security Association	In the cases of SR OS 2 SA per SC, each with a different SAK, each SC comprises a succession of SAs. Each SA is identified by the SC identifier, concatenated with a two-bit association number. The Secure Association Identifier (SAI), thus created, allows the receiving SecY to identify the SA, and thus the SAK used to decrypt and authenticate the received frame. The AN, and hence the SAI, is only unique for the SAs that can be used or recorded by participating SecYs at any instant. MACsec Key Agreement is responsible for creating and distributing SAKs to each of the SecYs in a CA. This key creation and distribution is independent of the cryptographic operation of each of the SecYs. The decision to replace one SA with its successor is made by the SecY that transmits using the SC, after MACsec Key Agreement has informed it that all the other SecYs are prepared to receive using that SA. No notification, other than receipt of a secured frame with a different SAI is sent to the receiver. A SecY must always be capable of storing SAKs for two SAs for each inbound SC, and of swapping from one SA to another without notice. Certain LAN technologies can reorder frames of different priority, so reception of frames on a single SC can use interleaved SA.
SAK: Security Association Key	SAK is the encryption key used to encrypt the datapath of MACsec.

2.3.2.3.5 MACsec Static CAK


MACsec uses SAs for encryption of packets. SA is a security relationship that provides security guarantees for frames transmitted from one member of a CA to the others. Each SA contains a single secret key (SAK) where the cryptographic operations used to encrypt the datapath PDUs.

SAK is the secret key used by an SA to encrypt the channel.

When enabled, MACsec uses a static CAK security mode. Two security keys, a connectivity association key (CAK) that secures control plane traffic and a randomly-generated secure association key (SAK) that secures data plane traffic are used to secure the point-to-point or point-to-multipoint Ethernet link. Both keys are regularly exchanged between both devices on each end of the Ethernet link to ensure link security.

Figure 8 illustrates MACsec generating the CAK.

Figure 8 MACsec Generating the CAK



The node initially needs to secure the control plain communication to distribute the SAKs between two or more members of a CA domain.


The securing of control plain is done via CAK. To generate the CAK, there are two main methods:

- EAPoL (SR OS does not support EAPoL)
- pre shared key (CAK and CKN values are configured manually via CLI). The following CAK and CKN rules apply.
 - CAK is a 32 hexadecimal characters for 128-bit key and 64 hexadecimal characters for 256-bit key depending on which algorithm is used for control plain encryption (for example, **aes-128-cmac** or **aes-256-cmac**).
 - CKN is a 32 octets char (64 hex) and it is the connectivity association key name which identifies the CAK. This allows each of the MKA participants to select which CAK to use to process a received MKPDU. MKA places no restriction on the format of the CKN, except that it must comprise an integral number of octets, between 1 and 32 (inclusive), and that all potential members of the CA use the same CKN.

- CKN and CAK must match on peers to create a MACsec Secure CA.

Figure 9 illustrates the MACsec control plane authentication and encryption.

Figure 9 MACsec Control Plane



sw0122

Once the CAK is generated, it can obtain two other keys. These keys are:

- KEK (Key Encryption Key) — used to wrap and encrypt the SAKs
- ICK (Integrity Connection Value (ICV) Key) — used to for integrity check of each MKPDU send between two CA

The key server then creates a SAK, that is shared with the CAs of the security domain, and that SAK secures all data traffic traversing the link. The key server will continue to periodically create and share a randomly-created SAK over the point-to-point link for as long as MACsec is enabled.

The SAK will be encrypted via the AES-CMAC, the KEK as encryption key, and ICK as integration key.

2.3.2.3.6 SAK Rollover

SR OS re-generates the SAK after the following events:

- when a new host has joined the CA domain and MKA hellos are received from this host
- when the sliding window is reaching the end of its 32-bit or 64-bit length
- when a new PSK is configured and a rollover of PSK has been executed

2.3.2.3.7 MKA

Each MACsec peer operates the MACsec Key Agreement Protocol (MKA). Each node can operate multiple MKAs base on the number of CA that it belongs to. Each instance of MKA is protected by a distinct secure connectivity Association key (CAK), that allows each PAE to ensure that information for a given MKA instance is only accepted from other peer that also possess that CAK, and therefore identifying themselves as members or potential members of the same CA. See [MACsec Static CAK](#) for a description of how the CAK identification is done via CKN.

MKA PDU Generation

[Table 9](#) describes the MKA PDUs generated for different traffic encapsulation matches.

Table 9 MKA PDU Generation

Configuration	Configuration Example (<s-tag>.<c-tag>)	MKA Packet Generation	Traffic pattern match/behavior
All-encap	Config>port>ethernet>do t1x>macsec sub-port 10 encap-match all-encap ca-name 10	untagged MKA packet	Matches all traffic on port, including untagged, single-tag, and double-tag. Default behavior; only available behavior in releases prior to 16.0.
UN-TAG	Config>port>ethernet>do t1x>macsec sub-port 1 encap-match untagged ca-name 2	untagged MKA packet	Matches only untagged traffic on port
802.1Q single S-TAG (specific S-TAG)	Config>port>ethernet>do t1x>macsec sub-port 2 encap-match dot1q 1 ca-name 3	MKA packet generated with S-TAG=1	Matches only single-tag traffic on port with tagID of 1

Table 9 MKA PDU Generation (Continued)

Configuration	Configuration Example (<s-tag>.<c-tag>)	MKA Packet Generation	Traffic pattern match/ behavior
802.1Q single S-TAG (any S-TAG)	Config>port>ethernet>do t1x>macsec sub-port 3 encap-match dot1q * ca-name 4	untagged MKA packet	Matches any dot1q single-tag traffic on port
802.1ad double tag (both tag have specific TAGs)	Config>port>ethernet>do t1x>macsec sub-port 4 encap-match qinq 1.1 ca-name 5	MKA packet generated with S-tag=1 and C- TAG=1	Matches only double-tag traffic on port with service tag of 1 and customer tag of 1
802.1ad double tag (specific S-TAG, any C- TAG)	Config>port>ethernet>do t1x>macsec sub-port 6 encap-match qinq 1.* ca-name 7	MKA packet generated with S-TAG=1	Matches only double-tag traffic on port with service tag of 1 and customer tag of any
802.1ad double tag (any S-TAG, any C-TAG)	Config>port>ethernet>do t1x>macsec sub-port 7 encap-match qinq *.* ca-name 8	untagged MKA packet	Matches any double-tag traffic on port

Tags in Clear Behavior by Traffic Encapsulation Types

[Table 10](#) describes how single or double tags in clear configuration under a connectivity association affects different traffic flow encryptions.

By default all tags are encrypted in CA. An MKA can be generated without any tags (un-tag) but the data being matched can be based on dot1q or q-in-q.

Table 10 **Tags in Clear Behavior**

Configuration	Traffic pattern match/behavior	Sub-port's CA configuration: No tag in clear text	Sub-port's CA configuration: Single-tag in clear text	Sub-port's CA configuration: Double tag in clear text
PORT All-encap	Matches all traffic on port, including untagged, single-tag, double-tag (Release 15.0 default behavior)	MKA PDU: untagged Untagged traffic: encrypted Single-tag traffic: encrypted, no tag in clear Double-tag traffic: encrypted, no tag in clear	MKA PDU: untagged Untagged traffic: in clear Single-tag traffic: encrypted, single-tag in clear Double-tag traffic: encrypted, single-tag in clear	MKA PDU: untagged Untagged traffic: in clear Single-tag traffic: in clear Double-tag traffic: encrypted, double-tag in clear
untagged	Matches only untagged traffic on port	MKA PDU: untagged Untagged traffic: encrypted Single-tag traffic: not matched by this MACsec policy Double-tag traffic: not matched by this MACsec policy	N/A	N/A
802.1Q single tag (specific tag)	Matches only single-tag traffic on port with the configured tag value	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: tag is encrypted Double-tag traffic: not matched by this MACsec policy	MKA PDU: same tag as the one configured under encap-match Untagged traffic: not matched by this MACsec policy Single-tag traffic: tag will be in clear Double-tag traffic: not matched by this MACsec policy	N/A

Table 10 **Tags in Clear Behavior (Continued)**

Configuration	Traffic pattern match/behavior	Sub-port's CA configuration: No tag in clear text	Sub-port's CA configuration: Single-tag in clear text	Sub-port's CA configuration: Double tag in clear text
802.1Q single tag (any tag)	Matches all single-tag traffic on port	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: encrypted Double-tag traffic: not matched by this MACsec policy	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: encrypted with single tag in clear Double-tag traffic: not matched by this MACsec policy	N/A
802.1ad double tag (both tag have specific values)	Matches only double-tag traffic on port with both configured tag values	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic matching both configured tags: encrypted, no tag in clear	MKA PDU: single tag, equal to S-TAG Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic matching both configured tags: single S-TAG in clear	MKA PDU: double tag, equal to the values configured under the encaps-match Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic matching both configured tags: encrypted, both tags in clear

Table 10 **Tags in Clear Behavior (Continued)**

Configuration	Traffic pattern match/behavior	Sub-port's CA configuration: No tag in clear text	Sub-port's CA configuration: Single-tag in clear text	Sub-port's CA configuration: Double tag in clear text
802.1ad double tag (specific S-TAG, any C-TAG)	Matches only double-tag traffic on port with the configured S-TAG	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic matching the configured S-TAG: encrypted, no tag in clear	MKA PDU: single tag, equal to S-TAG Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic matching the configured S-TAG: S-TAG tag in clear	MKA PDU: single tag, equal to S-TAG Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic matching the configured S-TAG: both tags in clear
802.1ad double tag (any S-TAG, any C-TAG)	Matches all double-tag traffic on port	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic: encrypted, no tag in clear	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic: S-TAG tag in clear	MKA PDU: untagged Untagged traffic: not matched by this MACsec policy Single-tag traffic: not matched by this MACsec policy Double-tag traffic: both tags in clear

2.3.2.3.8 Pre-shared Key

A peer may support the use of one or more pre-shared key (PSKs). An instance of MKA operates for each PSK that is administratively configured as active.

A pre-shared key may be created by NSP, or entered in CLI manually.

Each PSK is configured with two fields. The two fields are:

- CKN (connectivity association name)
- CAK value

The CAK name (CKN) is required to be unique per port among the configured sub-ports, and can be used to identify the key in subsequent management operations.

Each static CAK configuration can have two pre-shared key entries for rollover. The active PSK index dictates which CAK is being used for encrypting the MKA PDUs.

NSP has additional functionality to roll over and configure the PSK. The rollover via NSP can be based on a configured timer.

2.3.2.3.9 MKA Hello Timer

MKA uses a member identifier (MI) to identify each node in the CA domain.

A participant proves liveness to each of its peers by including their MI, together with an acceptably-recent message number (MN), in an MKPDU.

To avoid a new participant having to respond to each MKPDU from each partner as it is received, or trying to delay its reply until it is likely that MI MN tuples have been received from all potential partners, each participant maintains and advertises both a live peers list and a potential peers list. The live peers list includes all the peers that have included the participant's MI and a recent MN in a recent MKPDU. The potential peers list includes all the other peers that have transmitted an MKPDU that has been directly received by the participant or that were included in the Live Peers List of a MKPDU transmitted by a peer that has proved liveness. Peers are removed from each list when an interval of between MKA life time and MKA life time plus MKA Hello Time has elapsed since the participant's recent MN was transmitted. This time is sufficient to ensure that two or more MKPDUs will have been lost or delayed prior to the incorrect removal of a live peer.



Note: The specified use of the live and potential peer lists permits rapid removal of participants that are no longer active or attached to the LAN while reducing the number of MKPDUs transmitted during group formation. For example, a new participant will be admitted to an established group after receiving, then transmitting, one MKPDU.

[Table 11](#) shows the MKA participant timer values used on SR OS.

Table 11 MKA Participant Timer Values

Timer Use	Timeout (parameter)	Timeout (parameter)
Per participant periodic transmission, initialized on each transmission on expiry	MKA Hello Time or MKA Bounded Hello Time	2.0 0.5
Per peer lifetime, initialized when adding to or refreshing the Potential Peers List or Live Peers List, expiry causes removal from the list	MKA Life Time	6.0
Participant lifetime, initialized when participant created or following receipt of an MKPDU, expiry causes participant to be deleted		
Delay after last distributing a SAK, before the Key Server will distribute a fresh SAK following a change in the Live Peer List while the Potential Peer List is still not empty		

2.3.2.3.10 MACsec Capability, Desire, and Encryption Offset

802.1x-2010 had identified two fields in the MKA PDU. Those fields are:

- MACsec Capability
- Desire

MACsec Capability signals whether MACsec is capable of integrity and confidentiality. [Table 12](#) describes the four basic settings for MACsec Capability.

Table 12 MACsec Basic Settings

Setting	Description
0	MACsec is not implemented
1	Integrity without confidentiality

Table 12 MACsec Basic Settings (Continued)

Setting	Description
2	The following are supported: • Integrity without confidentiality • Integrity and confidentiality with a confidentiality offset of 0
3 ¹	The following are supported: • Integrity without confidentiality • Integrity and confidentiality with a confidentiality offset of 0, 30, or 5

Note:

1. SR OS supports (3) Integrity without confidentiality and Integrity and confidentiality with a confidentiality offset of 0, 30, or 50.

Encryption offset of 0, 30, or 50 starts from the byte after the SecTAG (802.1ae header). Ideally, the encryption offset should be configured for IPv4 (offset 30) and IPv6 (offset 50) to leave the IP header in the clear text. This will allow routers and switches to use the IP header for LAG or ECMP hashing.

2.3.2.3.11 Key Server

The participants, in a given MKA instance agree on a Key Server, are responsible for the following:

- deciding on the use of MACsec
- cipher suite selection
- SAK generation and distribution
- SA assignment
- identifying the CA when two or more CAs merge

Each participant in an MKA instance uses the Key Server priority (an 8-bit integer) encoded in each MKPDU to agree on the Key Server. Each participant selects the live participant advertising the highest priority as its Key Server whenever the Live Peers List changes, provided that highest priority participant has not selected another as its Key Server or is unwilling to act as the Key Server. If a Key Server cannot be selected, SAKs are not distributed. In the event of a tie for highest priority Key Server, the member with the highest priority SCI is chosen. For consistency with other uses of the SCI's MAC address component as a priority, numerically lower values of the Key Server Priority and SCI are accorded the highest priority.



Note: With SCI, each SC is identified by an SCI that comprises a globally unique MAC address and a Port Identifier, unique within the system that has been allocated that address.

2.3.2.3.12 SA Limits and Network Design

Each MACsec device supports 64 TX-SAs and 64 RX-SAs. An SA (Security Association) is the key to encrypt or decrypt the data.

As per 802.1AE document, each SecY contains a SC. An SC is a unidirectional concept; for example, Rx-SC or Tx-SC. Each SC contains at least one SA for encryption on Tx-SC and decryption on Rx-SC. Also, for extra security, each SC should be able to roll over the SA and, as such, it is recommended for each SC to have two SAs for rollover purposes.

MACsec phy will be known as a MACsec security zone. Each MACsec security zone supports 64Tx-SAs and 64 RX-SAs. Assuming two SAs per SC for SA rollover, then each zone will support 32 RX-SC and 32 TX-SC.

[Table 13](#) describes the port mapping to security zones.


Table 13 Port Mapping to Security Zone

MDA	Ports in Security Zone 1	Ports in Security Zone 2	Ports in Security Zone 3	SA Limit per Security Zone
12-port SFP+/SFP MDA-e	Ports 1, 2, 3, 4	Ports 5, 6, 7, 8	Ports 9, 10, 11, 12	Rx-SA = 64 Tx-SA = 64

2.3.2.3.13 P2P (Switch to Switch) Topology


In a point-to-point topology, each router needs a single security zone and single Tx-SC for encryption and a single Rx-SC for decryption. Each SC has two SAs. In total for point-to-point topology, four SAs are needed, two RxSA for RxSC1 and two TXSA for TxSC1. See [Figure 10](#).

Figure 10 Switch Point to Switch Point Topology



2.3.2.3.14 P2MP (Switch to Switch) Topology

In a multipoint topology with N nodes, each node will need a single TxSC and N RxSC, one for each one of the peers. As such, 64 max RX-SA per security zone will translate to 32 Rx-SCs, which will break down to only 32 peers (for example, only 33 nodes in the multipoint topology per security zone, from each node perspective there is one TxSC and 32 RxSC).

Figure 11 Switch Multi-point to Switch Multi-point Topology

In [Figure 11](#), when the 34th node joins the multi-point topology, all other 33 nodes already part of this domain will not have any SAs to create an RxSC for this 34th node. However, the 34th node will have a TxSC and will accept 32 peers. The 34th node will start to transmit and encrypt the PDUs based on its TxSC. However, because all other nodes do not have a SC for this SAI, they will drop all Rx PDUs.

It is recommended to ensure that a multicast domain, for a single security zone, does not exceed 32 peers or the summation of all the nodes, in a security zone's CA domain, do not exceed 33. This is the same is if a security zone has four CAs, the summation of all nodes in the four CAs should be 33 or less.

2.3.2.3.15 SA Exhaustion Behavior

In [SA Limits and Network Design](#), it was explained that the a security zone has 64 RxSAs and 64 TxSAs. Two RxSAs will be used for each RxSC for rollover purposes and two TxSAs will be used for TxSC for rollover purposes. This translates to 32 peers per security zone.

Under each port, a max-peer parameter can be configured. This parameter assigns how many peers are allowed on that port.



Note: The operator must ensure that the number of peers do not exceed the limit of maximum peers per security zone or maximum peers per port (for example, exceeds the **port max-peer** parameter).

If the maximum peer is exceeded, the peer connectivity will be random in case of a node failure or packet loss. Peers will join the CA randomly, on a first-come first-served basis.



Caution: It is highly recommended that the operator ensures the maximum peer does not exceed per-security zone or port.

2.3.2.3.16 Clear Tag Mode

In most Layer 2 networks, MAC forwarding is done via destination MAC address. The 802.1AE standard dictates that any field after source and destination MAC address and after the SecTAG is required to be encrypted. This includes the 802.1Q tags. In some VLAN switching networks, it might be desired to leave the 802.1Q tag in clear text.

SR OS allows the configuration of 802.1Q tag, in clear text, by placing the 802.1Q tag before the SecTAG or encrypted, by placing it after the SecTAG.

[Table 14](#) lists the MACsec encryption of 802.1Q tags when the clear-tag is configured on SR OS.

Table 14 MACsec Encryption of 802.1Q Tags with Clear-tag Configured

Unencrypted format	Clear-tag-mode configuration	Pre-encryption (Tx)	Pre-decryption (Rx)
Single tag (dot1q)	Single-tag	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTAG
Single tag (dot1q)	Double-tag	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTAG
Double tag (q-in-q)	Single-tag	DA, SA, TPID1, VID1, IPID2, VID2, Etype	DA, SA, TPID1, VID1, SecTAG
Double tag (q-in-q)	Double-tag	DA, SA, TPID1, VID1, IPID2, VID2, Etype	DA, SA, TPID1, VID1, IPID2, VID2, SecTAG

2.3.2.3.17 802.1X Tunneling and Multihop MACsec

MACsec is an Ethernet packet and, as with any other Ethernet packet, can be forwarded through multiple switches via Layer 2 forwarding. The encryption and decryption of the packets will be performed via the 802.1x (MKA) capable ports.

To ensure that MKA is not terminated on any intermediate switch or router, the user can enable 802.1x tunneling on the corresponding port.

A sample check to see if tunneling is enabled, is provided below.

```
*A:SwSim28>config>port>ethernet>dot1x# info
-----
tunneling
```

By enabling tunneling, the 802.1X MKA packets will be transit that port, without being terminated, as such MKA negotiation will not happen on a port that has 802.1X tunneling enabled.

2.3.2.3.18 EAPoL Destination Address

The MKA packets are transported over EAPoL with a multicast destination MAC address. At some point, it might be desired to have the MKA have a point-to-point connection to a peer node over a Layer 2 multihop cloud. In this case, the EAPoL destination MAC address can be set to the peer MAC address. This will force the MKA to traverse multiple nodes and establish an MKA session with the specific peer.

2.3.2.3.19 Mirroring Consideration

Mirroring is performed prior to the MACsec encryption engine. Therefore, if a port is MACsec-enabled and also, that same port is being mirrored, all the mirrored packets will be in clear text.

2.3.2.4 SONET/SDH Port Attributes

One OC-3/STM-1 port is supported on the CMA. One OC-3/STM-1 port is supported on the MDA. OC-3/STM-1 ports are also supported on TDM satellites. The ports can be configured for either SONET or SDH operation. SONET ports are configured for channelized OC-3 operation. SDH ports can be configured for channelized STM-1 operation.

The port's transmit clock rate can be node or loop timed. The port's receive clock rate can be used as a synchronization source for the system. The Section Trace (C1) byte can be configured by the user to ensure proper physical cabling. The port can activate and deactivate local line and internal loopbacks.

All SONET/SDH line alarms are configurable to be either enabled (default) or disabled. Link hold timers can be configured in 100ms increments to control link up and link down indications. The line signal degradation bit error rate (ber-sd) threshold and the line signal failure bit error rate (ber-sf) threshold can be configured.

The CMAs, MDAs, and TDM satellite support all standard SR OC-3/STM-1 SFP optics including multi-mode, intermediate reach, and long reach. Single fiber mode is not supported.

The CMA contains 3 LEDs for power, status and link state of port #1. The MDA contains LEDs for power, status and one for each link state. The power LED is blue if power is connected and off if no power is present. The status LED is green when operationally up, amber when operationally down, off when administratively shutdown and blinking green during initialization. The link state LED is green when the link is established; amber when the link is down; and unlit when the port is shutdown.

When an Ethernet port is configured in WAN mode (xgig wan), you can change certain SONET/SDH parameters to reflect the SONET/SDH requirements for this port. See [SONET/SDH Port Commands](#) for more information.

2.3.2.5 SONET/SDH Path Attributes

Any CES path can only be configured to operate in access mode. Each path has a configurable text description. The SONET/SDH signal label byte (C2) is configurable. The SONET/SDH path trace string (J1) is configurable. Payload scrambling can not be enabled on CES paths. The valid SONET and SDH path configurations are shown in [Table 15](#).

Table 15 Valid SONET and SDH Path Configurations

Framing	Path Configuration Options Per Physical Port	Max Number of Paths Per Physical Port	Max Number of Paths Per TDM Satellite Port
SDH	STM1>AUG1>VC4>TUG3>TUG2>VC12>E1 STM1>AUG1>VC3>TUG2>VC12>E1	63 E1 or 512 n*64 kb/s	63 E1
SONET	OC3>STS1 SPE>DS3>E1		—

Table 15 Valid SONET and SDH Path Configurations (Continued)

Framing	Path Configuration Options Per Physical Port	Max Number of Paths Per Physical Port	Max Number of Paths Per TDM Satellite Port
SONET	OC3>STS1 SPE>VT GROUP>VT1.5 SPE>DS1	84 DS1 or 512 n*64 kb/s	84 DS1
SONET	OC3>STS1 SPE>DS3	3 DS3	—
SONET	OC3>STS1 SPE>DS3>DS1	84 DS1, 63 E1 or 512 n*64 kb/s	—
SDH	STM1>AUG1>VC4>TUG3>TUG2>TU11>VC11>DS1 STM1>AUG1>VC3>TUG2>VC11>DS1	84 DS1 or 512 n*64 kb/s	84 DS1
SDH	STM1>AUG1>VC3>DS3>DS1	84 DS1, 63 E1 or 512 n*64 kb/s	—
SDH	STM1>AUG1>VC4>TUG3>VC3>E3 STM1>AUG1>VC3>E3	3 E3	—
SDH	STM1>AUG1>VC3>DS3	3 DS3	—
SDH	STM1>AUG1>VC3>DS3>E1	3 DS3	—

All SONET/SDH path alarms are configurable to be either enabled (the default) or disabled. The MTU size is configurable per path in the range of 512 to 2092. The path uses a default MTU size set to equal the largest possible CES packet size.

Load balancing options are not applicable to channelized CES paths.

When an Ethernet port is configured in WAN mode (xgig wan), you can change certain SONET/SDH parameters to reflect the SONET/SDH requirements for this port. See [SONET/SDH Path Commands](#) for details.

2.3.2.6 Multilink Frame Relay

MLFR is a bundling capability allowing users to spray FR frame fragments over multiple T1/E1 links. This allows a dynamic provisioning of additional bandwidth by adding incremental bandwidth between T1/E1 and DS3/E3. A MLFR bundle increases fault tolerance and improves QoS characteristics since one single large frame of low priority cannot block a higher priority frame.

A MLFR supports up to eight (8) member links and a maximum of 128 bundles with up to 336 T1/252 E1 members links can be configured per MDA. NxDS0 circuits or higher speed circuits are not supported.

The MLFR implementation supports FRF.16.1 bundle link integrity protocol to verify serviceability of a member link.

2.3.2.6.1 MLFR Bundle Data Plane

FRF.16.1 reuses the UNI/NNI fragmentation procedures defined in FRF.12. Frames on all FR SAP on the MLFR bundle have the UNI/NNI fragmentation header added regardless if they are fragmented or not. A separate sequence number state machine is used for each FR SAP configured on the bundle. The fragmentation threshold is configurable in the range 128 to 512 bytes.

In order to provide priority based scheduling of the FR SAP fragments over the bundle links, the user configures a FR scheduling class for each FR SAP configured on the bundle. As in MC-MLPPP, four scheduling classes are supported.

A separate fragmentation context is used by each FR SAP. FR SAPs of the same scheduling class share the same egress FR scheduling class queue with fragments of each SAP packets stored contiguously. The fragments from each scheduling class queue are then sprayed over the member links. Furthermore, the user may select the option to not fragment but spray the FR frames with the fragmentation header included over the member links.

Received fragments over the member links are re-assembled on a per SAP basis to re-create the original FR frame.

A user is not allowed to add an FR SAP with FRF.12 e2e fragmentation enabled to an MLFR bundle. Conversely, the user cannot enable FRF.12 e2e fragmentation on an FR SAP configured on an MLFR bundle. If an FR frame with the e2e fragmentation header is received on a bundle, it is forwarded if the FR SAP is part of an Fpipe service. It will be discarded if the FR SAP is part of any other service.

Note that the operator must disable LMI before adding a link to an MLFR bundle. Also, the operator must shut down the bundle in order to change the value of the fragmentation threshold.

An FR SAP configured on an MLFR bundle can be part of a VLL, VPLS, IES, or VPRN service.

2.3.2.6.2 MLFR Bundle Link Integrity Protocol

FRF.16.1 defines a MLFR Bundle Link Integrity Protocol which verifies the serviceability of a member link. If a problem is found on the member link the link integrity protocol will identify the problem, flag the link as unusable, and adjust the Bundle's available bandwidth. For MLFR Bundles the link integrity protocol is always enabled.

For each member link of a bundle the link integrity protocol will do the following:

- Confirm frame processing capabilities of each member link.
- Verify membership of a link to a specific remote bundle.
- Report to the remote end of the member link the bundle to which the link belongs
- Detect loopbacks on the member link. This is always enabled on the 7750 SR. The near-end monitors the magic number Information Element (IE) sent by the far-end and if its value matches the one it transmitted in ten consecutive control messages, it sends a `remove_link` message to the far-end and brings the link down. The near-end will attempt to add the link until it succeeds.
- Estimate propagation delay on the member link. The differential delay is calculated as follows in the 7750 SR implementation. Every time the near-end sends an `add_link` or Hello message to the far-end, it includes the Timestamp Information Element (IE) with the local time the packet was sent. FRF16.1 standard requires that the remote equipment includes the timestamp IE and copies the received timestamp value unchanged if the sender included this IE. When the far-end node sends back the ACK for these messages, the near-end calculates the round trip time. The 7750 SR implementation maintains a history of the last "N" round-trip-times that were received. It takes the fastest of these samples for each member link to find out the member link with the fastest RTT. Then for each link it calculates the difference between the fastest links RTT, and the RTT for the current link. The user has the option to coordinate link removal between the local and remote equipment. Note, however, that in the SR 7750 implementation, the addition of a link will be hitless but the removing a link is not.

Specifically, the MLFR Bundle Link Integrity Protocol defines the following control messages:

- `ADD_LINK`
- `ADD_LINK_ACK`
- `ADD_LINK_REJ`
- `HELLO`
- `HELLO_ACK`
- `REMOVE_LINK`
- `REMOVE_LINK_ACK`

The control messages are encapsulated in a single-fragment frame where the C-bit, the B-bit, and the E-bit are all set. The details of the message format are given in FRF.16.1. [Table 16](#) lists the user configured control parameters with values as specified in FRF.16.1.

Table 16 FRF.16.1 Values

Parameter	Default Value	Minimum Value	Maximum Value
Timer T_HELLO	10 seconds	1 second	180 seconds
Timer T_ACK	4 seconds	1 second	10
Count N_MAX_RETRY	2	1	5

T_HELLO Timer - this timer controls the rate at which hello messages are sent. Following a period of T_HELLO duration, a HELLO message is transmitted onto the Bundle Link.

Note that T_HELLO Timer is also used, during the Bundle Link adding process, as an additional delay before re-sending an ADD_LINK message to the peer Bundle Link when this peer Bundle Link does not answer as expected.

T_ACK Timer - this timer defines the maximum period to wait for a response to any message sent onto the Bundle Link before attempting to retransmit a message onto the Bundle Link.

N_RETRY - this counter specifies the number of times a retransmission onto a Bundle Link will be attempted before an error is declared and the appropriate action taken.

2.3.2.7 FRF.12 End-to-End Fragmentation

The user enables FRF.12 e2e fragmentation on a per FR SAP basis. A fragmentation header is added between the standard Q.922 header and the payload. This header consists of a 2-byte Network Layer Protocol ID (NLPID) of value 0xB1 to indicate e2e fragmentation payload and a 2-byte containing the Beginning bit (B-bit), the End-bit (E-bit), the Control bit (C-bit), and the Sequence Number field.

The following is the mode of operation for the fragmentation in the transmit direction of the FR SAP. Frames of all the FR SAP forwarding class queues are subject to fragmentation. The fragmentation header is, however, not included when the frame size is smaller than the user configured fragmentation size. The SAP transmits all fragments of a frame before sending the next full or fragmented frame. The fragmentation threshold is configurable in the range 128 to 512 bytes. In the receive direction, the SAP accepts a full frame interleaved with fragments of another frame to interoperate with other vendor implementations.

An FR SAP with FRF.12 e2e fragmentation enabled can be part of a VPLS service, an IES service, a VPRN service, an Ethernet VLL service, or an IP VLL service. This SAP cannot be part of a FR VLL service or an FRF.5 VLL service. However, fragmented frames received on such VLLs will be passed transparently as in current implementation.

2.3.2.7.1 SAP Fragment Interleaving Option

This option provides a different mode of operation for the fragmentation in the transmit direction of the FR SAP than in the default behavior of a FRF.12 end-to-end fragmentation. It allows for the interleaving of high-priority frames and fragments of low-priority frames.

When the interleave option is enabled, only frames of the FR SAP non expedited forwarding class queues are subject to fragmentation. The frames of the FR SAP expedited queues are interleaved, with no fragmentation header, among the fragmented frames. In effect, this provides a behavior like in MLPPP Link Fragment Interleaving (LFI). The receive direction of the FR SAP supports both modes of operation concurrently, for example, with and without fragment interleaving.

2.3.2.8 FRF.12 UNI/NNI Link Fragmentation

The user enables FRF.12 UNI/NNI link fragmentation on a per FR circuit basis. All FR SAPs configured on this circuit are subject to fragmentation. A fragmentation header is added on top of the standard Q.922 header. This header consists of 2 bytes containing the beginning bit (B-bit), the End-bit (E-bit), the Control bit (C-bit), and the sequence number field. The fragmentation header is included on frames of all SAPs regardless if the frame size is larger or not than the fragment size.

The FECN, BECN, and DE bits of all fragments of a given FR frame are set to the same value as the original frame. The FECN, BECN, and DE bits of a re-assembled frame are set to the logical OR of the corresponding bits on the constituent fragments.

The operator must delete all configured FR SAPs on a port before enabling or disabling FRF.12 UNI/NNI on that port. Also, the user must shut down the port in order to change the value of the fragmentation threshold.

A FR SAP on a FR circuit with FRF.12 UNI/NNI fragmentation enabled can be part of a VLL, VPLS, IES, or VPRN service.

QoS for a link with FRF.12 UNI/NNI fragmentation is the same as for a MLFR bundle. The FR class queue parameters and its scheduling parameters are configured by applying an egress QoS profile to an FRF.12 UNI/NNI port. The FR scheduling class ingress re-assembly timeout is not applicable to a FRF.12 UNI/NNI port.

2.3.2.9 MLFR/FRF.12 Support of APS, BFD, and Mirroring Features

The following APS support is provided:

- Single-chassis APS is supported on a SONET/SDH port with FRF.12 UNI/NNI fragmentation enabled on the port or on a constituent TDM circuit.
- Single-chassis APS is supported on a SONET/SDH port with FRF.12 e2e fragmentation enabled on one or more FR SAPs on the port or on a constituent TDM circuit.
- Single-chassis APS is not supported on a SONET/SDH port with MLFR bundles configured.
- Multi-chassis APS is not supported on a SONET/SDH port with FR encapsulation configured on the port or on a constituent TDM circuit.
- APS is not supported on TDM satellite.

The following BFD support is provided:

- BFD is supported on an IP interface configured over a FR SAP with e2e fragmentation enabled.
- BFD is supported on an IP interface configured over a FR SAP on a port or channel with UNI/NNI fragmentation enabled.
- BFD is not supported on an FR SAP configured on an MLFR bundle.

The following mirroring support is provided:

- Port mirroring and FR SAP mirroring on an MLFR bundle.
- IP mirroring for an FR SAP on an MLFR bundle.
- A mirror source can be an MLFR bundle or a FR SAP on an FR bundle.
- Mirror destinations must be FR SAPs and must not be part of an APS group or an MLFR bundle.

2.3.2.10 Multilink Point-to-Point Protocol (MLPPP)

Multilink point-to-point protocol is defined in the IETF RFC 1990, *The PPP Multilink Protocol (MP)*, and provides a way to distribute data across multiple links within an MLPPP bundle to achieve high bandwidth. MLPPP allows for a single frame to be fragmented and transmitted across multiple links. This allows for lower latency and also allows for a higher maximum receive unit (MRU).

MP is negotiated during the initial LCP option negotiations of a standard PPP session. A router indicates to its peer that it is willing to perform MLPPP by sending the MP option as part of the initial LCP option negotiation. This negotiation indicates the following:

1. The system offering the option is capable of combining multiple physical links into one logical link;
2. The system is capable of receiving upper layer protocol data units (PDU) fragmented using the MP header and reassembling the fragments back into the original PDU for processing;
3. The system is capable of receiving PDUs of size N octets where N is specified as part of the option even if N is larger than the maximum receive unit (MRU) for a single physical link.

Once MLPPP has been successfully negotiated, the sending system is free to send PDUs encapsulated and/or fragmented with the MP header.

MP introduces a new protocol type with a protocol ID (PID) of 0x003d. [Figure 12](#) and [Figure 13](#) show the MLPPP fragment frame structure. Framing to indicate the beginning and end of the encapsulation is the same as that used by PPP, and described in PPP in HDLC-like framing [RFC 1662]. MP frames use the same HDLC address and control pair value as PPP, namely: Address - 0xFF and Control - 0x03. The two octet protocol field is also structured the same as in PPP encapsulation.

Figure 12 MLPPP 24-bit Fragment Format





Figure 13 MLPPP 12-bit Fragment Format



The required and default format for MP is the 24-bit format. During the LCP state the 12-bit format can be negotiated. The SR-series routers can support and negotiate the alternate 12-bit frame format.

2.3.2.10.1 Protocol Field (PID)

The protocol field is two octets its value identifies the datagram encapsulated in the Information field of the packet. In the case of MP the PID also identifies the presence of a 4-octet MP header (or 2-octet, if negotiated).

A PID of 0x003d identifies the packet as MP data with an MP header.

The LCP packets and protocol states of the MLPPP session follow those defined by PPP in RFC 1661, *The Point-to-Point Protocol (PPP)*. The options used during the LCP state for creating an MLPPP NCP session are described below.

2.3.2.10.2 B & E Bits

The B&E bits are used to indicate the epoch of a packet. Ingress packets to the MLPPP process will have an MTU, which may or may not be larger than the MRRU of the MLPPP network. The B&E bits manage the fragmentation of ingress packets when it exceeds the MRRU.

The B-bit indicates the first (or beginning) packet of a given fragment. The E-bit indicates the last (or ending) packet of a fragment. If there is no fragmentation of the ingress packet both B&E bits are set true (=1).

2.3.2.10.3 Sequence Number

Sequence numbers can be either 12 or 24 bits long. The sequence number is zero for the first fragment on a newly constructed AVC bundle and increments by one for each fragment sent on that bundle. The receiver keeps track of the incoming sequence numbers on each link in a bundle and reconstructs the desired unbundled flow through processing of the received sequence numbers and B&E bits. For a detailed description of the algorithm refer to RFC 1990.

2.3.2.10.4 Information Field

The Information field is zero or more octets. The Information field contains the datagram for the protocol specified in the protocol field.

The MRRU will have the same default value as the MTU for PPP. The MRRU is always negotiated during LCP.

2.3.2.10.5 Padding

On transmission, the Information field of the ending fragment may be padded with an arbitrary number of octets up to the MRRU. It is the responsibility of each protocol to distinguish padding octets from real information. Padding must not be added to any but the last fragment (the E-bit set true).

2.3.2.10.6 FCS

The FCS field of each MP packet is inherited from the normal framing mechanism from the member link on which the packet is transmitted. There is no separate FCS applied to the reconstituted packet as a whole if transmitted in more than one fragment.

2.3.2.10.7 LCP

The Link Control Protocol (LCP) establishes the connection through an exchange of configure packets. This exchange is complete, and the LCP opened state entered, once a Configure-Ack packet has been both sent and received.

LCP allows for the negotiation of multiple options in a PPP session. MLPPP is somewhat different than PPP and therefore the following options are set for MLPPP and not negotiated:

- No async control character map
- No link quality monitoring
- No compound frames
- No self-describing-padding

Any non-LCP packets received during this phase must be silently discarded.

2.3.2.10.8 Link Fragmentation and Interleaving Support

Link Fragmentation and Interleaving (LFI) provides the ability to interleave high priority traffic within a stream of fragmented lower priority traffic. This feature helps avoid excessive delays to high priority, delay-sensitive traffic over a low-speed link. This can occur if this traffic type shares a link with lower priority traffic that utilizes much larger frames. Without this ability, higher priority traffic must wait for the entire packet to be transmitted before being transmitted, which could result in a delay that is too large for the application to function properly.

For example, if VoIP traffic is being sent over a DS-1 or fractional DS-1 which is also used for Best Effort Internet traffic, LFI could be used so the small (usually 64-128B) VoIP packets can be transmitted between the transmission of fragments from the lower priority traffic.

Figure 14 shows the sequence of events as low priority and high priority frames arrive and are handled by LFI.

Figure 14 Frame Sequence of Events




Fig_2

1. A low priority frame arrives in the low priority queue. At this particular instant, there are no packets in the high priority queue so low priority frame is de-queued and passed to the fragmentation mechanism for MLPPP.
2. The original packet is divided into 'n' fragments based on the size of the packet and the fragment threshold configuration.

3. The fragments are then transmitted out the egress port.
4. After the transmission of the fragments has begun, high priority frames arrive in the high priority queue.
5. The transmission of the remaining fragments stops and the high priority packets are transmitted out the egress interface. Note that high priority packets are not fragmented.
6. When the high priority traffic is transmitted, the remaining lower priority fragments are then transmitted.

On the ingress side, LFI requires that the ingress port can receive non-fragmented packets within the fragment stream and pass these packets directly on to the forwarding engine and then continue with the reassembly process for the fragmented frames.



2.3.2.11 Multi-Class MLPPP

Multi-class MLPPP (MC-MLPPP) allows for the prioritization of multiple types of traffic flowing between the cell site routers and the mobile operator's aggregation routers. MC-MLPPP is an extension of the MLPPP standard which allows multiple classes of service to be transmitted over a MLPPP bundle. Originally, link fragmentation and interleaving (LFI) was added to MLPPP that allowed two classes, but in some applications, two classes of service can be insufficient.

The MLPPP header includes two class bits to allow for up to four classes of service. This enhancement to the MLPPP header format is detailed in RFC 2686, *The Multi-Class Extension to Multi-Link PPP*. This allows multiple classes of services over a single MLPPP connection and allows the highest priority traffic to be transmitted over the MLPPP bundle with minimal delay regardless of the order in which packets are received.

[Table 17](#) shows the header formats and [Table 18](#) shows the original MLPP header format and the enhanced header format.

Table 17 Header Formats

Original MLPPP Header Format	MC-MLPPP Short Sequence Header Format
	

The new MC-MLPPP header format uses the two (previously unused) bits before the sequence number as the class identifier. This allows four distinct classes of service to be identified into separate re-assembly contexts.

2.3.2.11.1 QoS in MC-MLPPP

If the user enables the multiclass option under an MLPPP bundle, the MDA egress data path provides a queue for each of the four classes of MLPPP. The user configures the required number of MLPPP classes to use on a bundle. The forwarding class of the packet, as determined by the ingress QoS classification, determines the MLPPP class for the packet and hence which of the four egress MDA queues to store the packet. The mapping of forwarding class to MLPPP class is a function of the user configurable number of MLPPP classes. The default mapping for a 4-class, 3-class, and 2-class MLPPP bundle is shown in [Table 18](#).

Table 18 Default Packet Forwarding Class to MLPPP Class Mapping

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class 4-class bundle	MLPPP Class 3-class bundle	MLPPP Class 2-class bundle
7	NC	Expedited	0	0	0
6	H1	Expedited	0	0	0
5	EF	Expedited	1	1	1

Table 18 Default Packet Forwarding Class to MLPPP Class Mapping (Continued)

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class 4-class bundle	MLPPP Class 3-class bundle	MLPPP Class 2-class bundle
4	H2	Expedited	1	1	1
3	L1	Non-Expedited	2	2	1
2	AF	Non-Expedited	2	2	1
1	L2	Non-Expedited	3	2	1
0	BE	Non-Expedited	3	2	1

Table 19 shows a different mapping enabled when the user applies one of three pre-defined egress QoS profiles in the 4-class bundle configuration only.

Table 19 Packet Forwarding Class to MLPPP Class Mapping

FC ID	FC Name	Scheduling Priority (Default)	MLPPP Class (MLPPP Egress QoS profile 1, 2, and 3)
7	NC	Expedited	0
6	H1	Expedited	0
5	EF	Expedited	1
4	H2	Expedited	2
3	L1	Non-Expedited	2
2	AF	Non-Expedited	2
1	L2	Non-Expedited	2
0	BE	Non-Expedited	3


The MLPPP class queue parameters and its scheduling parameters are also configured by applying one of the three pre-defined egress QoS profiles to an MLPPP bundle.

Table 20 and Figure 15 provide the details of the class queue threshold parameters. Packets marked with a high drop precedence, such as out-of-profile, by the service or network ingress QoS policy will be discarded when any class queue reaches the OOP threshold. Packet with a low drop precedence marking, such as in-profile, will be discarded when any class queue reaches the max threshold.

Table 20 MLPPP Class Queue Threshold Parameters

	Class 0		Class 1		Class 2		Class 3	
Queue Threshold (in ms @ Available bundle rate)	Max	Oop	Max	Oop	Max	Oop	Max	Oop
2-Class Bundle Default Egress QoS Profile	250	125	750	375	N/A	N/A	N/A	N/A
3-Class Bundle Default Egress QoS Profile	50	25	200	100	750	375	N/A	N/A
4-Class Bundle Default Egress QoS Profile	10	5	50	25	150	75	750	375
4-Class Bundle Egress QoS Profile 1	25	12	5	3	200	100	1000	500
4-Class Bundle Egress QoS Profile 2	25	12	5	3	200	100	1000	500
4-Class Bundle Egress QoS Profile 3	25	12	5	3	200	100	1000	500

Figure 15 MLPPP Class Queue Thresholds for In-Profile and Out-of-Profile Packets



QSSG258


Table 21 and Figure 16 provide the details of the class queue scheduling parameters.

Table 21 MLPPP Class Queue Scheduling Parameters

		WRR Parameters		
4-class MLPPP Egress QoS Profile	MIR	W1	W2	W3
Profile 1	85%	<1%	66%	33%

Table 21 MLPPP Class Queue Scheduling Parameters (Continued)

		WRR Parameters		
Profile 2	90%	<1%	89%	10%
Profile 3	85%	<1%	87%	12%

Figure 16 MLPPP Class Queue Scheduling Scheme

Note that all queue threshold and queue scheduling parameters are adjusted to the available bundle rate. If a member link goes down or a new member link is added to the bundle, the scheduling parameters MIR, W1, W2, W3, as well as the per class queue thresholds OOP and max are automatically adjusted to maintain the same values.

Class 0 queue is serviced at MLPPP at available bundle rate. Class 1 queue is guaranteed a minimum service rate but is allowed to share additional bandwidth with class 2 and 3 queues based on the configuration of WRR weight W1.

Class queues 2 and 3 can be given bandwidth guarantee by limiting MIR of class 1 queue to less than 100% and by setting the WRR weights W1, W2, and W3 to achieve the desired bandwidth distribution among all three class queues.

Note that there is one queue per bundle member link to carry link control packets, such as LCP: PPP, and which are serviced with strict priority over the 4 class queues (not shown).

In the default 2-class, 3-class, and 4-class egress QoS profile, the class queues are service with strict priority in ascending order of class number.

Ingress MLPPP Class Reassembly

For an MLPPP bundle with the multi-class option enabled, there is a default profile for setting the re-assembly timer value for each class. When the pre-defined MLPPP ingress QoS profile 1 is applied to a 4-class bundle, the values of the timers are modified as shown in [Table 22](#).

Table 22 MLPPP Ingress QoS Profile: Reassembly Timers (msec)

	Class 0	Class 1	Class 2	Class 4
MLPPP ingress QoS default profile (2-Class bundle)	25ms	25ms	NA	NA
MLPPP ingress QoS default profile (3-Class bundle)	25ms	25ms	25ms	NA
MLPPP ingress QoS default profile (4-Class bundle)	25ms	25ms	100ms	1000ms
MLPPP ingress QoS profile 1 (4-class bundle)	10	10	100	1000

Configuring MC-MLPPP QoS Parameters

A 4-class MLPPP bundle can be configured with user-defined MLPPP QoS attributes. This feature cannot be used with MC-MLPPP bundles with fewer than 4 classes or with non-multiclass bundles.

The following describe the parameters and the configuration processes and rules

1. The user creates an ingress QoS profile in the **mlppp-profile-ingress** context, to configure a preferred value of the ingress per-class re-assembly timer. Ingress QoS profile 1 is reserved for the pre-defined profile with parameter values shown in [Table 22](#). The user is allowed to edit this profile and change the parameter values. When a user creates a profile with a profile-id greater than 1, or performs the no option command on the parameter, the parameter's default value will always be the 1 in [Table 22](#) for ingress QoS Profile #1 regardless of the parameter value the edited Profile 1 has at that point.
2. The user creates an egress QoS profile in the **mlppp-profile-egress** context to configure preferred values for the per-class queue and queue scheduling parameters. The user can also configure system forwarding class mapping to the MLPPP classes. Egress QoS profiles 1, 2, and 3, are reserved for the pre-defined profiles with parameter values shown in [Table 19](#), [Table 20](#), or [Table 21](#). Users can edit these profiles and change the parameter values. When a user creates a profile with a profile-id higher than 3, or when the user specifies the no

- option command on the parameter, the default value will be the one shown in [Table 19](#), [Table 20](#), or [Table 21](#) for the egress QoS Profile 1. This is regardless of the parameter value the edited profiles have at that point in time.
3. A maximum of 128 ingress and 128 egress QoS profiles can be created on the system.
 4. The values of the ingress per-class re-assembly timer are configured in the ingress QoS profile.
 5. The mapping of the system forwarding classes to the MLPPP Classes are configured in the egress QoS profile. There is a many-to-one relationship between the system FC and an MLPPP class. See [Table 19](#) for the mapping when one of the three pre-defined 4-class egress QoS profiles is selected.
 6. The maximum size for each MLPPP class queue in units of msec at the available bundle rate is configured in the egress QoS profile. This is referred to as max in [Figure 15](#) and as max-queue-size in CLI. The out-of-profile threshold for an MLPPP class queue, referred to as oop in [Figure 15](#), is not directly configurable and is set to 50% of the maximum queue size rounded up to the nearest higher integer value.
 7. The MLPPP class queue scheduling parameters is configured in the egress QoS profile. The minimum information rate, referred to as **MIR** in [Figure 16](#) and **mir** in CLI, applies to Class 1 queue only. The MIR parameter value is entered as a percentage of the available bundle rate. The WRR weight, referred to as W1, W2, and W3 in [Figure 16](#) and weight in CLI, applies to class 1, class 2, and class 3 queues. Note that W1 in [Figure 16](#) is not configurable and is internally set to a value of 1 such that Class 1 queue shares 1% of the available bundle rate when the sum of W1, W2, and W3 equals 100. W2 and W3 weights are integer values and are user configurable such that Class 2 queue shares $(W2/(W1 + W2 + W3))$ and Class 3 queue shares $(W3/(W1 + W2 + W3))$ of the available bundle rate.
 8. The user applies the ingress and egress QoS profiles to a 4-class MLPPP bundle for the configured QoS parameter values to take effect on the bundle.
 9. The following operations require the bundles associated with a QoS profile to be shutdown to take effect.
 - A change of the numbered ingress or egress QoS profile associated with a bundle.
 - A change of the bundle associated ingress or egress QoS profile from default profile to a numbered profile and vice-versa.
 10. The following operations can be performed without shutting down the associated bundles:
 - Changes to any parameters in the ingress and egress QoS profiles.

The CLI commands for the creation of ingress and egress QoS profiles and configuration of the individual QoS parameters are described in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Quality of Service Guide*.

2.3.2.12 Cisco HDLC

Cisco HDLC (cHDLC) is an encapsulation protocol for information transfer. It is a bit-oriented synchronous data-link layer protocol that specifies a data encapsulation method on synchronous serial links using frame characters and checksums.

cHDLC monitors line status on a serial interface by exchanging keepalive request messages with peer network devices. It also allows routers to discover IP addresses of neighbors by exchanging Serial Link Address Resolution Protocol (SLARP) (see [SLARP](#)) address-request and address-response messages with peer network devices.

The basic frame structure of a cHDLC frame is shown in [Table 23](#). This frame structure is similar to PPP in an HDLC-link frame (RFC 1662, *PPP in HDLC-like Framing*). The differences to PPP in and HDLC-like frames are in the values used in the address, control, and protocol fields.

Table 23 cHDLC I-Frame

Flag	Address	Control	Protocol	Information Field	FCS
0x7E	0x0F/0x8F	0x00	—	—	16/32 bits

- Address field — The values of the address field include: 0x0F (unicast), 0x8F (broadcast).
- Control field — The control field is always set to value 0x00.
- Protocol field — The following values are supported for the protocol field:

[Table 24](#) shows the cHDLC protocol fields.

Table 24 cHDLC Protocol Fields

Protocol	Field Value
IP	0x0800
Cisco SLARP	0x8035
ISO CLNP/ISO ES-IS DSAP/SSAP1	0xFEFE

- Information field — The length of the information field is in the range of 0 to 9 kbytes.

- FCS field — The FCS field can assume a 16-bit or 32-bit value. The default is 16-bits for ports with a speed equal to or lower than OC-3, and 32-bits for all other ports. The FCS for cHDLC is calculated in the same manner and same polynomial as PPP.

2.3.2.12.1 SLARP

A cHDLC interface on a Nokia router will transmit a SLARP address resolution reply packet in response to a received SLARP address resolution request packet from peers. The cHDLC interface will not transmit SLARP address resolution request packets.

For the SLARP keepalive protocol, each system sends the other a keepalive packet at a user-configurable interval. The default interval is 10 seconds. Both systems must use the same interval to ensure reliable operation. Each system assigns sequence numbers to the keepalive packets it sends, starting with zero, independent of the other system. These sequence numbers are included in the keepalive packets sent to the other system. Also included in each keepalive packet is the sequence number of the last keepalive packet received from the other system, as assigned by the other system. This number is called the returned sequence number. Each system keeps track of the last returned sequence number it has received. Immediately before sending a keepalive packet, it compares the sequence number of the packet it is about to send with the returned sequence number in the last keepalive packet it has received. If the two differ by 3 or more, it considers the line to have failed, and will not route higher-level data across it until an acceptable keepalive response is received.

There is interaction between the SLARP address resolution protocol and the SLARP keepalive protocol. When one end of a serial line receives a SLARP address resolution request packet, it assumes that the other end has restarted its serial interface and resets its keepalive sequence numbers. In addition to responding to the address resolution request, it will act as if the other end had sent it a keepalive packet with a sequence number of zero, and a returned sequence number the same as the returned sequence number of the last real keepalive packet it received from the other end.

2.3.2.12.2 SONET/SDH Scrambling and C2-Byte

SONET/SDH scrambling and overhead for cHDLC follow the same rules used for POS (RFC 2615, *PPP over SONET/SDH*).

The two key SONET/SDH parameters are scrambling and signal-label (C2-byte). Scrambling is off by default. The default value of the C2-byte is 0xCF. These two parameters can be modified using the CLI. The other SONET overhead values (for example, j0) follow the same rules as the current POS implementation.

SONET/SDH scrambling and overhead for cHDLC is not supported on TDM satellite.


2.3.2.12.3 Timers

Cisco HDLC (cHDLC) has two timers associated with the protocol, the keepalive interval and the timeout interval. The keepalive interval sends periodic keepalive packets. The receiver process expects to receive a keepalive packet at the rate specified by the keepalive interval. The link is declared down if the receiver process does not receive a keepalive within the timeout interval. The link is declared up when the number of continual keepalive packets received equals the up-count.

It is recommended that the nodes at the two endpoints of the cHDLC link are provisioned with the same values.

2.3.2.13 Automatic Protection Switching (APS)

APS is designed to protect SONET/SDH equipment from linear unidirectional or bidirectional failures. The Network Elements (NEs) in a SONET/SDH network constantly monitor the health of the network. When a failure is detected, the network proceeds through a coordinated pre-defined sequence of steps to transfer (or switchover) live traffic to the backup facility (protection facility). This happens very quickly to minimize lost traffic. Traffic remains on the protection facility until the primary facility (working facility) fault is cleared, at which time the traffic may optionally be reverted to the working facility. An example is shown in [Figure 17](#).

Figure 17 APS Protection (Single Chassis APS) and Switchover

Note that “facility” in the router’s context refers to the physical line (including intermediate transport/switching equipment) and directly attached line terminating hardware (SFP module, MDA and IOM). “Circuit” is also a term used for a link/facility (working-circuit).

A 1+1 APS group contains two circuits.

APS is configured on a port by port basis. If all ports on an MDA or IOM need to be protected then each port on the MDA or IOM must be individually added into an APS group.

Working and protection circuits can be connected to a variety of types of network elements (ADMs, DACSes, ATM switches, routers) and serve as an access or network port providing one or more services or network interfaces to the router. APS-protected SONET/SDH ports may be further channelized, and may contain bundled channels MLPPP or IMA Bundle Protection Groups). The ports may be one of a variety of encapsulation types as supported by the MDA including PPP, ATM, FR and more. For information about MDAs, port types, switching modes, bundles and encapsulations supported with APS, see [APS Applicability, Restrictions and Interactions](#).

This section discusses the different APS architectures and their implementations.

- [Single Chassis and Multi-Chassis APS](#)
- [APS Switching Modes](#)
- [APS Channel and SONET Header K Bytes](#)

- [Revertive Switching](#)
- [Bidirectional 1+1 Switchover Operation Example](#)
- [Protection of Upper Layer Protocols and Services](#)
- [APS User-Initiated Requests](#)
- [APS and SNMP](#)
- [APS Applicability, Restrictions and Interactions](#)
- [Sample APS Applications](#)

2.3.2.13.1 Single Chassis and Multi-Chassis APS

APS can operate in a single chassis configuration (SC-APS) or in a multi-chassis configuration (MC-APS).

An SC-APS group can span multiple ports, MDAs or IOMs within a single node whereas as MC-APS can span two separate nodes as shown in [Table 25](#).

Table 25 SC-APS versus MC-APS Protection

	Single Chassis APS	Multi-Chassis APS
Short form name	SC-APS	MC-APS
Link failure protection (including intermediate transmission equipment failure)	Yes	Yes
Optical/electrical module (SPF, XFP) failure protection	Yes	Yes
MDA failure protection	Yes	Yes
IOM failure protection	Yes	Yes
Node failure protection	No	Yes

The support of SC-APS and MC-APS depends on switching modes, MDAs, port types and encaps. For a definitive description of the MDAs, port types, switching modes, bundles and encapsulations supported with APS, see [APS Applicability, Restrictions and Interactions](#).

APS on a Single Node (SC-APS)


In a single chassis APS both circuits of an APS group are terminated on the same node.

The working and protect lines of a single chassis APS group can be:

- Two ports on the same MDA
- Two ports on different MDAs but on the same IOM
- Two ports on different MDAs on two different IOMs (installed in different slots)
- Two ports on two TDM satellites

If the working and protection circuits are on the same MDA, protection is limited to the physical port and the media connecting the two devices. If the working and protection circuits are on different IOMs then protection extends to MDA or IOM failure. [Figure 18](#) shows a configuration that provides protection against circuit, port, MDA or IOM failure on the 7750 SR connected to an Add-Drop-Multiplexer (ADM).

Figure 18 SC-APS Group with MDA and IOM Protection



APS Across Two Nodes (MC-APS)

Multi-Chassis APS functionality extends the protection offered by SC-APS to include protection against nodal (7750 SR) failure by configuring the working circuit of an APS group on one 7750 SR node while configuring the protect circuit of the same APS group on a different 7750 SR node.

These two nodes connect to each other with an IP link to establish an MC-APS signaling path between the two 7750 SRs. Note that the working circuit and the protect circuit must have compatible configurations (such as the same speed, framing, and port-type). The relevant APS groups in both the working and protection routers must have same group ID, but they can have different names (for example, group port descriptions). Although the working and protection routers can be different platforms (7750 SR-7 and a 7750 SR-c12), switchover performance may be

impacted so it is recommended to avoid a mix of platforms in the same MC-APS group where possible. The configuration consistency between the working circuit/router and the protection circuit/router is not enforced by the 7750 SR. Service or network-specific configuration data is not signaled nor synchronized between the two service routers.


Signaling is provided using the direct connection between the two service routers. A heartbeat protocol can be used to add robustness to the interaction between the two routers. Signaling functionality includes support for:

- APS group matches between service routers.
- Verification that one side is configured as a working circuit and the other side is configured as the protect circuit. In case of a mismatch, a trap (incompatible neighbor) is generated.
- Change in working circuit status is sent from the working router to keep the protect router in sync.
- Protect router, based on K1/K2 byte data, member circuit status, and external request, selects the active circuit, and informs the working router to activate or de-activate the working circuit.

Note that external requests like lockout, force, and manual switches are allowed only on the APS group having the protection circuit.

The [Figure 19](#) shows a Multi-Chassis APS group being used to protect against link, port, MDA, IOM or node failure.

Figure 19 MC-APS Group Protects Against Node Failure



OSSG657

2.3.2.13.2 APS Switching Modes

APS behavior and operation differs based on the switching mode configured for the APS group as shown in [Table 26](#). Several switching modes are supported in the router.

The switching mode affects how the two directions of a link behave during failure scenarios and how APS tx operates.

Unidirectional / Bidirectional configuration must be the same at both sides of the APS group. The APS protocol (K byte messages) exchange switching mode information to ensure that both nodes can detect a configuration mismatch.

- If one end of an APS group is configured in a Unidirectional mode (Uni 1+1 Sig APS or Uni 1+1 Sig+Data APS) then the other end must also be configured in a Unidirectional mode (Uni 1+1 Sig+Data APS).
- If one end of an APS group is configured in a Bidirectional mode then the other end must also be configured in Bidirectional mode.

Table 26 APS Switching Modes

	Bidirectional 1+1 Signaling APS	Unidirectional 1+1 Signaling APS	Unidirectional 1+1 Signaling and Datapath APS
Short form name	Bidir 1+1 Sig APS	Uni 1+1 Sig APS	Uni 1+1 Sig+Data APS
CLI	bi-directional	uni-directional	uni-1plus1
Interworks with a standards compliant APS implementation	Yes	Yes	Yes
Full 1+1 APS standards-based signaling	Yes	Yes	Yes
Data is transmitted simultaneously on both links/ circuits (1+1 Data)	No	No	Yes

The support of switching modes depends on SC-APS / MC-APS, MDAs, port types and encaps. For a definitive description of the MDAs, port types, switching modes, bundles and encapsulations supported with APS, see [APS Applicability, Restrictions and Interactions](#).

Bidirectional 1+1 Signaling APS

In Bidir 1+1 Sig APS switching mode the Tx data is sent on the active link only (it is not bridged to both links simultaneously). 1+1 signaling, however, is used for full interoperability with signaling-compliant 1+1 architectures.

In the ingress direction (Rx), the decision to accept data from either the working or protection circuit is based on both locally detected failures/degradation and on what circuit the far-end is listening on (as indicated in the K bytes). If the far-end indicates that it has switched its active receiver, then the local node will also switch its receiver (and Tx) to match the far-end. If the local Rx changes from one circuit to another it notifies the far end using the K bytes.

In the egress direction (Tx), the data is only transmitted on the active circuit. If the active Rx changes, then Tx will also change to the same circuit.

Bidirectional 1+1 Signaling APS ensures that both directions of active data flow (including both Rx) are using the same link/circuit (using the two directions of the same fiber pair) as required by the APS standards. If one end of the APS group changes the active receiver, it will signal the far end using the K bytes. The far end will then also change its receiver to listen on the same circuit.

Because the router transmits on active circuits only and keeps active TX and RX on the same port, both local and remote switches are required to restore the service.

The APS channel (bytes K1 and K2 in the SONET header – K bytes) exchanges requests and acknowledgments for protection switch actions. In Bidirectional 1+1 Signaling APS switching mode, the router sends correct status on the K bytes and requires the far-end to also correctly update/send the K-bytes to ensure that data is transmitted on the circuit on which the far-end has selected as its active receiver.

Line alarms are processed and generated independently on each physical circuit.

In Bidirectional 1+1 Signaling APS mode, the highest priority local request is compared to the remote request (received from the far end node using an APS command in the K bytes), and whichever has the greater priority is selected. The relative priority of all events that affect APS 1+1 protection is listed in the [Table 27](#) in descending order. The requests can be automatically initiated (such as signal failure or signal degrade), external (such as lockout, forced switch, request switch), and state requests (such as revert-time timers, and so on).

Unidirectional 1+1 Signaling APS

In Uni 1+1 Sig APS switching mode the Tx data is sent on the active link only (it is not bridged to both links simultaneously). 1+1 signaling, however, is used for full interoperability with signaling-compliant 1+1 architectures.

In the ingress direction (Rx), the decision to accept data from either the working or protection circuit is based on both locally detected failures/degradation and on what circuit the far-end is listening on (as indicated in the K bytes). Although it is not required in the APS standards, the system's implementation of Unidirectional 1+1 Signaling APS uses standards based signaling to keep both the Rx and Tx on the same circuit / port. If the far-end indicates that it has switched its active receiver, then the local node will also switch its receiver (and Tx) to match the far-end. If the local Rx changes from one circuit to another it notifies the far end using the K bytes.

In the egress direction (Tx), the data is only transmitted on the active circuit. If the active Rx changes, then Tx will also change to the same circuit.

Because the router transmits on active circuits only and keeps active TX and RX on the same port, both local and remote switches are required to restore the service. For a single failure a data outage is limited to a maximum of 100 milliseconds.

The APS channel (bytes K1 and K2 in the SONET header – K bytes) exchanges requests and acknowledgments for protection switch actions. In Unidirectional 1+1 Signaling APS switching mode, the router sends correct status on the K bytes and requires the far-end to also correctly update/send the K-bytes to ensure that data is transmitted on the circuit on which the far-end has selected as its active receiver.

Line alarms are processed and generated independently on each physical circuit.

In Unidirectional 1+1 Signaling APS switching mode:

- K-bytes are generated/transmitted based on local request/condition only (as required by the APS signaling).
- Local request priority is compliant to 1+1 U-APS specification.
- RX and TX are always forced on to the same (active) circuit (bi-directional). This has the following caveats:
 - If an APS switch is performed due to a local condition, then the TX direction will be moved as well to the newly selected RX circuit (old inactive). The router will send LAIS on the old active TX circuit to force the remote end to APS switch to the newly active circuit. Note that some local request may not cause an APS switch when a remote condition prevents both RX and TX direction to be on the same circuit (for example an SD detected locally on a working circuit will not cause a switch if the protection circuit is locked out by the remote end).

- If the remote end indicates an APS switch and the router can RX and TX on the circuit newly selected by the remote end, then the router will move its TX direction and will perform an APS switch of its RX direction (unless the router already TX and RX on the newly selected circuit).
- If the remote end indicates an APS switch and the router cannot RX and TX on the circuit newly selected by the remote end (for example due to a higher priority local request, like a force request or manual request, and so on), then L-AIS are sent on the circuit newly selected by the remote end to force it back to the previously active circuit.
- The sent L-AIS in the above cases can be either momentary or persistent. The persistent L-AIS is sent under the following conditions:
 - On the protection circuit when the protection circuit is inactive and cannot be selected due to local SF or Lockout Request.
 - On the working circuit as long as the working circuit remains inactive due to a local condition. The persistent L-AIS is sent to prevent revertive switching at the other end.

In all other cases a momentary L-AIS is sent. The system provides debugging information that informs operators about the APS-induced L-AIS.

Unidirectional 1+1 Signaling and Datapath APS

Uni 1+1 Sig+Data APS supports unidirectional switching operations, 1+1 signaling and 1+1 data path.

In the ingress direction (Rx) switching is done based on local requests only as per the APS specifications. K-bytes are used to signal the far end the APS actions taken.

In the egress direction (Tx), the data is transmitted on both active and protecting circuits.

Each end of the APS group may be actively listening on a different circuit.

The APS channel (bytes K1 and K2 in the SONET header) exchanges APS protocol messages.

In Uni 1+1 Sig+Data APS a received L-RDI signal on the active circuit does not cause that circuit (port) to be placed out of service. The APS group can continue to use that circuit as the active receiver. This behavior is not configurable.

Uni 1+1 Sig+Data APS also supports configurable:

- Debounce timers for signal failure and degradation conditions
- Suppression of L-RDI alarm generation

2.3.2.13.3 APS Channel and SONET Header K Bytes

The APS channel (bytes K1 and K2 in the SONET header) exchanges APS protocol messages for all APS modes.

K1 Byte

The switch priority of a request is assigned as indicated by bits 1 through 4 of the K1 byte (as described in the rfc3498 APS-MIB); see [Table 27](#).

Table 27 K1 Byte, Bits 1 to 4: Type of Request

Bit 1234	Condition
1111	Lockout of protection
1110	Force switch
1101	SF - High priority
1100	SF - Low priority
1011	SD - High priority
1010	SD - Low priority
1001	(not used)
1000	Manual switch
0111	(not used)
0110	Wait-to-restore
0101	(not used)
0100	Exercise
0011	(not used)
0010	Reverse request
0001	Do not revert
0000	No request

The channel requesting switch action is assigned by bits 5 through 8. When channel number 0 is selected, the condition bits show the received protection channel status. When channel number 1 is selected, the condition bits show the received working channel status. Channel values of 0 and 1 are supported.

[Table 28](#) shows bits 5 to 8 of a K1 byte and K2 Bits 1 to 4 and the channel number code assignments.

Table 28 K1 Byte, Bits 5 to 8 (and K2 Bits 1 to 4), Channel Number Code Assignments

Channel Number Code	Channel and Notes
0	Null channel. SD and SF requests apply to conditions detected on the protection line. For 1+1 systems, Forced and Request Switch requests apply to the protection line (for the 7750 SR only). Only code 0 is used with Lockout of Protection request.
1 to 14	Working channel. Only code 1 applies in a 1+1 architecture. Codes 1 through n apply in a 1:n architecture (for the 7750 SR only). SD and SF conditions apply to the corresponding working lines.
15	Extra traffic channel. May exist only when provisioned in a 1:n architecture. Only No Request is used with code 15.

K2 Byte

The K2 byte indicates the bridging actions performed at the line-terminating equipment (LTE), the provisioned architecture and mode of operation.

The bit assignment for the K2 byte is listed in [Table 29](#).

Table 29 K2 Byte Functions

Bits 1 to 8	Function
1 to 4	Channel number. The 7750 SR supports only values of 0 and 1.
5	0 Provisioned for 1+1 mode 1 Provisioned for 1:n mode

Table 29 K2 Byte Functions (Continued)

Bits 1 to 8	Function
6 to 8	111 Line AIS
	110 Line RDI
	101 Provisioned for bi-directional switching
	100 Provisioned for uni-directional switching
	011 (reserved for future use)
	010 (reserved for future use)
	001 (reserved for future use)
	000 (reserved for future use)

Differences in SONET/SDH Standards for K Bytes

SONET and SDH standards are slightly different with respect to the behavior of K1 and K2 Bytes.

[Table 30](#) shows the differences between the two standards.

Table 30 Differences Between SONET and SDH Standards

	SONET	SDH	Comments
SONET/SDH standards use different codes in the transmitted K1 byte (bits 1-4) to notify the far-end of a signal fail/signal degrade detection.	1100 for signal fail 1010 for signal degrade 1101 unused 1011 unused	1101 for signal fail 1011 for signal degrade 1100 unused 1010 unused	None.
SONET systems signal the switching mode in bits 5-8 of the K2 byte whereas SDH systems do not signal at all.	101 for bi-dir 100 for uni-dir	Not used. 000 is signaled in bits 5 to 8 of K2 byte for both bi-directional as well as uni-directional switching.	SONET systems raise a mode mismatch alarm as soon as a mismatch in the TX and RX K2 byte (bits 5 to 8) is detected. SDH systems do not raise the mode mismatch alarm.

Failures Indicated by K Bytes

The following sections describe failures indicated by K bytes.

APS Protection Switching Byte Failure

An APS Protection Switching Byte (APS-PSB) failure indicates that the received K1 byte is either invalid or inconsistent. An invalid code defect occurs if the same K1 value is received for 3 consecutive frames (depending on the interface type (framer) used, the 7750 SR may not be able to strictly enforce the 3 frame check per GR-253 and G.783/G.841) and it is either an unused code or irrelevant for the specific switching operation. An inconsistent APS byte defect occurs when no three consecutive received K1 bytes of the last 12 frames are the same.

If the failure detected persists for 2.5 seconds, a Protection Switching Byte alarm is raised. When the failure is absent for 10 seconds, the alarm is cleared. This alarm can only be raised by the active port operating in bi-directional mode.

APS Channel Mismatch Failure

An APS channel mismatch failure (APS-CM) identifies that there is a channel mismatch between the transmitted K1 and the received K2 bytes. A defect is declared when the received K2 channel number differs from the transmitted K1 channel number for more than 50 ms after three identical K1 bytes are sent. The monitoring for this condition is continuous, not just when the transmitted value of K1 changes.

If the failure detected persists for 2.5 seconds, a channel mismatch failure alarm is raised. When the failure is absent for 10 seconds, the alarm is cleared. This alarm can only be raised by the active port operating in a bi-directional mode.

APS Mode Mismatch Failure

An APS mode mismatch failure (APS-MM) can occur for two reasons. The first is if the received K2 byte indicates that 1:N protection switching is being used by the far-end of the OC-N line, while the near end uses 1+1 protection switching. The second is if the received K2 byte indicates that uni-directional mode is being used by the far-end while the near-end uses bi-directional mode.

This defect is detected within 100 ms of receiving a K2 byte that indicates either of these conditions. If the failure detected persists for 2.5 seconds, a mode mismatch failure alarm is raised. However, it continues to monitor the received K2 byte, and should it ever indicate that the far-end has switched to a bi-directional mode the mode mismatch failure clearing process starts. When the failure is absent for 10 seconds, the alarm is cleared, and the configured mode of 1+1 bidirectional is used.

APS Far-End Protection Line Failure

An APS far-end protection line (APS-FEPL) failure corresponds to the receipt of a K1 byte in 3 consecutive frames that indicates a signal fail (SF) at the far end of the protection line. This forces the received signal to be selected from the working line.

If the failure detected persists for 2.5 seconds, a far-end protection line failure alarm is raised. When the failure is absent for 10 seconds, the alarm is cleared. This alarm can only be raised by the active port operating in a bi-directional mode.

2.3.2.13.4 Revertive Switching

The APS implementation also provides the revertive and non-revertive modes with non-revertive switching as the default option. In revertive switching, the activity is switched back to the working port after the working line has recovered from a failure (or the manual switch is cleared). In non-revertive switching, a switch to the protection line is maintained even after the working line has recovered from a failure (or if the manual switch is cleared).

A revert-time is defined for revertive switching so frequent automatic switches as a result of intermittent failures are prevented. A change in this value takes effect upon the next initiation of the wait to restore (WTR) timer. It does not modify the length of a WTR timer that has already been started. The WTR timer of a non-revertive switch can be assumed to be infinite.

In case of failure on both working and the protection line, the line that has less severe errors on the line will be active at any point in time. If there is signal degrade on both ports, the active port that failed last will stay active. When there is signal failure on both ports, the working port will always be active. The reason is that the signal failure on the protection line is of a higher priority than on the working line.

2.3.2.13.5 Bidirectional 1+1 Switchover Operation Example

[Table 31](#) outlines the steps that a bi-directional protection switching process will go through during a typical automatic switchover.

Table 31 Actions for the Bi-directional Protection Switching Process

Status	APS Commands Sent in K1 and K2 Bytes on Protection Line		Action	
	B -> A	A -> B	At Site B	At Site A
No failure (Protection line is not in use).	No request.	No request.	No action.	No action.
Working line Degraded in direction A->B.	SD on working channel 1.	No request.	Failure detected, notify A and switch to protection line.	No action.
Site A receives SD failure condition.	Same.	Reverse request.	No action.	Remote failure detected, acknowledge and switch to protection line.
Site B receives Reverse request.	Same.	Same.	No action.	No action.

2.3.2.13.6 Annex B (1+1 Optimized) Operation

Operation and behavior conform with Annex B of ITU.T G.841 can be configured for an APS group. Characteristics of this mode include are the following:

- Annex B operates in non-revertive bi-directional switching mode only as defined in G.841.
- Annex B operates with 1+1 signaling, but 1:1 data path where by data is transmitted on the active link only.
- K bytes are transmitted on both circuits.

Due to the request/reverse-request nature of an Annex B switchover, the data outage is longer than a typical (non Annex B single chassis) APS switchover. IMA bundles that are protected with Annex B APS have to resynchronize after a switchover. It is recommended to use maintenance commands (**tools>perform>aps...**) for planned switchovers (not MDA or IOM shutdown) to minimize the outage.

Annex B APS Outage Reduction Optimization

Typical standard Annex B behavior when a local SF is detected on the primary section (circuit), and this SF is the highest priority request on both the local side and from the remote side as per the APS specifications, is to send a request to the remote end and then wait until a reverse request is received before switching over to the secondary section. To reduce the recovery time for traffic, the router will switch over to the secondary section immediately upon detecting the local SF on the primary section instead of waiting for the reverse request from the remote side. If the remote request is not received after a period of time then an “PSB Failure is declared” event is raised (Protection Switching Byte Failure – indicates an inconsistent or invalid Rx K1 Bytes), and the APS group on the local side switches back to the primary section.

When the remote side is in Lockout, and a local SF is detected then a reverse request will not be received by the local side. In this case, the traffic will no longer flow on the APS group since neither the primary nor secondary sections can carry traffic, and the outage reduction optimization will cause a temporary switchover from the primary to the secondary and then back again (which causes no additional outage or traffic issue since neither section is usable). If this temporary switchover is not desired then it is recommended to either perform Lockout from the router side, or to Lockout from both sides, which will avoid the possibility of the temporary switchover.

Failures detected on the secondary section cause immediate switch over as per the Annex B specification. There is no outage reduction optimization in the router for this case as it is not needed.

Some examples of events that can cause a local SF to be detected include: a cable being cut, laser transmitter or receiver failure, a port administratively “shutdown”, MDA failure or shutdown, IOM failure or shutdown.



Note: In Annex B operation, all switch requests are for a switch from the primary section to the secondary section. Once a switch request clears normally, traffic is maintained on the section to which it was switched by making that section the primary section. The primary section may be working circuit 1 or working circuit 2 at any particular moment.

2.3.2.13.7 Protection of Upper Layer Protocols and Services

APS prevents upper layer protocols and services from being affected by the failure of the active circuit.

The following example with figures and description illustrate how services are protected during a single-chassis APS switchover.

Figure 20 shows an example in which the APS working circuit is connected to IOM-1/MDA-1 and the protection circuit is connected to IOM-2/MDA-1. In this example, assume that the working circuit is currently used to transmit and receive data.

Figure 20 APS Working and Protection Circuit Example




Fig. 4

Switchover Process for Transmitted Data

For packets arriving on all interfaces that need to be transmitted over APS protected interfaces, the next hop associated with all these interfaces are programmed in all Flexible Fast-Path complexes in each MDA with a logical next-hop index. This next hop-index identifies the actual next-hop information used to direct traffic to the APS working circuit on IOM-1/MDA-1.

All Flexible Fast-Path complexes in each MDA are also programmed with next hop information used to direct traffic to the APS protect circuit on IOM-2/MDA-1. When the transmitted data needs to be switched from the working to the protect circuit, only the relevant next hop indexes need to be changed to the pre-programmed next-hop information for the protect circuit on IOM-2/MDA-1.

Although the control CFM/CPM on the SF/CPM blade initiates the changeover between the working to protect circuit, the changeover is transparent to the upper layer protocols and service layers that the switchover occurs.

Physical link monitoring of the link is performed by the CPU on the relevant IOM for both working and protect circuits.

Switchover Process for Received Data

The Flexible Fast-Path complexes for both working and protect circuits are programmed to process ingress. The inactive (protect) circuit however is programmed to ignore all packet data. To perform the switchover from working circuit to the protect circuit the Flexible Fast-Path complex for the working circuit is set to ignore all data while the Flexible Fast-Path complex of the protect circuit will be changed to accept data.

The ADM or compatible head-end transmits a valid data signal to both the working and protection circuits. The signal on the protect line will be ignored until the working circuit fails or degrades to the degree that requires a switchover to the protect circuit. When the switchover occurs all services including all their QoS and filter policies are activated on the protection circuit.

2.3.2.13.8 APS User-Initiated Requests

The following subsections describe APS user-initiated requests.

Lockout Protection

The lockout of protection disables the use of the protection line. Since the **tools>perform>aps>lockout** command has the highest priority, a failed working line using the protection line is switched back to itself even if it is in a fault condition. No switches to the protection line are allowed when locked out.

Request Switch of Active to Protection

The request or manual switch of active to protection command switches the active line to use the protection line unless a request of equal or higher priority is already in effect. If the active line is already on the protection line, no action takes place.

Request Switch of Active to Working

The request or manual switch of active to working command switches the active line back from the protection line to the working line unless a request of equal or higher priority is already in effect. If the active line is already on the working line, no action takes place.

Forced Switching of Active to Protection

The forced switch of active to protection command switches the active line to the protection line unless a request of equal or higher priority is already in effect. When the forced switch of working to protection command is in effect, it may be overridden either by a lockout of protection or by detecting a signal failure on the protection line. If the active line is already on the protection line, no action takes place.

Forced Switch of Active to Working

The forced switch of active to working command switches the active line back from the protection line to the working unless a request of equal or higher priority is already in effect.

Exercise Command

The exercise command is only supported in the bi-directional mode of the 1+1 architecture. The exercise command is specified in the **tools>perform>aps>force>exercise** context and exercises the protection line by sending an exercise request over the protection line to the tail-end and expecting a reverse request response back. The switch is not actually completed during the exercise routine.

2.3.2.13.9 APS and SNMP

SNMP Management of APS uses the APS-MIB (from rfc3498) and the TIMETRA-APS-MIB.

[Table 32](#) shows the mapping between APS switching modes and MIB objects.

Table 32 Switching Mode to MIB Mapping

switching-mode	TIMETRA-APS-MIB tApsProtectionType	APS-MIB apsConfigDirection
Bidir 1+1 Sig APS (bi-directional)	onePlusOneSignalling (1)	bidirectional (2)
Uni 1+1 Sig APS (uni-directional)	onePlusOneSignalling (1)	unidirectional (1)

Table 32 Switching Mode to MIB Mapping (Continued)

switching-mode	TIMETRA-APS-MIB tApsProtectionType	APS-MIB apsConfigDirection
Uni 1+1 Sig+Data APS (uni-1plus1)	onePlusOne (2)	unidirectional (1)

apsConfigMode in the APS-MIB is set to onePlusOneOptimized for Annex B operation.

2.3.2.13.10 APS Applicability, Restrictions and Interactions



Note: The Release Notes for the relevant SR OS release should be consulted for details about APS restrictions.

Table 33 shows the supported APS mode combinations.

Table 33 Supported APS Mode Combinations

	Bidirectional 1+1 Signaling APS	Unidirectional 1+1 Signaling APS	Unidirectional 1+1 Signaling and Datapath APS
Single Chassis APS (SC-APS)	Supported ¹	Supported	Supported for 7750 SR-c4/12 platforms only
Multi-Chassis APS (MC-APS)	Supported	Not supported	Not supported

Note:

1. TDM satellite supports this mode only.

APS and Bundles

Bundles (such as IMA and MLPPP) can be protected with APS through the use of Bundle Protection Groups (BPGRP). For APS-protected bundles, all members of a working bundle must reside on the working port of an APS group. Similarly all members of a protecting bundle must reside on the protecting circuit of that APS group. Bundles are not supported on TDM satellite.

IMA APS protection is supported only when the router is connected to another piece of equipment (possibly through an ADM) running a single IMA instance at the far end. By design, the IMA APS implementation is expected to keep the IMA protocol up as long as the far end device can tolerate some frame loss. Similarly, the PPP protocol state machine for PPP channels and MLPPP bundles remains UP when a switchover occurs between the working and protect circuits.

When APS protects IMA groups, IMA control cells, but not user traffic, are sent on the inactive circuit (as well as the active) to keep the IMA protocol up during an APS switch.

For details on MLFR/FRF.12 support with APS see [MLFR/FRF.12 Support of APS, BFD, and Mirroring Features](#).

APS Switchover Impact on Statistics

All SAP-level statistics are retained with an APS switch. A SAP will reflect the data received regardless of the number of APS switches that has occurred. ATM statistics, however, are cleared after an APS switch. Thus, any ATM statistics viewed on an APS port are only the statistics since the current active member port became active.

Physical layer packet statistics on the APS group reflect what is currently on the active member port.

Port and path-level statistics follow the same behavior as described above.

Any SONET physical-layer statistics (for example, B1,B2,B3,...) on the APS port are only what is current on the active APS member port.

Supported APS MDA/Port Combinations

[Table 34](#) shows examples of the port types that can be paired to provide APS protection. Both ports must be the same type and must be configured at the same speed.

Table 34 MDA/Port Type Pairing for APS

MDA Type	Circuit Emulation (CES) For example: m4-choc3-ces-sfp	Channelized Any Service Any Port (ASAP) For example: m1-choc12-as-sfp	SONET/SDH Satellite
Circuit Emulation (CES) For example: m4-choc3-ces-sfp	Supported	—	—
Channelized Any Service Any Port (ASAP) For example: m1-choc12-as-sfp	—	Supported	—
SONET/SDH Satellite	—	—	Supported

APS Switchover During CFM/CPM Switchover

An APS switchover immediately before, during or immediately after a CFM/CPM switchover may cause a longer outage than normal.

Removing or Failure of a Protect MDA

The detection of a CMA/MDA removal or a CMA/MDA failure can take additional time. This can affect the APS switchover time upon the removal or failure of a protection CMA/MDA. If the removal is scheduled during maintenance, it is recommended that the port and/or protect circuit be shutdown first to initiate an APS switchover before the CMA/MDA maintenance is performed.

Mirroring Support

Mirroring parameters configured on a specific port or service, are maintained during an APS failover.

2.3.2.13.11 Sample APS Applications


The following subsections provide sample APS application examples.

Sample APS Application: MLPPP with SC-APS and MC-APS on Channelized Interfaces

The 7750 SR supports APS on channelized interfaces. This allows the router to be deployed as the radio access network (RAN) aggregation router which connects the base transceiver station (BTS) and the radio network controller (RNC).

[Figure 21](#) shows an example of MLPPP termination on APS protected channelized OC-n/STM-n links. This example illustrates the following:

- SC-APS (the APS circuits terminate on the same node aggregation router A).
- APS protecting MLPPP bundles (bundles are between the BTS and aggregation router A, but APS operates on the SONET links between the DACS and the aggregation router).
- APS on channelized access interfaces (OC-3/OC-12 links).

Figure 21 SC-APS MLPPP on Channelized Access Interfaces Example


OSSG142

Figure 22 shows an APS group between a digital access cross-connect system (DACS) and a pair of aggregation routers. At one end of the APS group both circuits (OC-3/STM-1 and/or OC-12/STM-4 links) are terminated on the DACS and at the other end each circuit is terminated on a different aggregation routers to provide protection against router failure. The MLPPP bundle operates between the BTS and the aggregation routers. At any one time only one of the two aggregation routers is actually terminating the MLPPP bundle (whichever aggregation router is processing the active APS circuit).

This example shows the following:

- MC-APS (the APS circuits terminate on different aggregation routers)
- APS protecting MLPPP bundles (bundles are between the BTS and the aggregation routers but APS operates on the SONET links between the DACS and the aggregation routers)
- APS on channelized access interfaces (OC-3/OC-12 links)


Figure 22 MC-APS MLPPP on Channelized Access Interfaces Example



OSSG143

Sample APS Application: MC-APS for ATM SAP with ATM VPLS Service

In [Figure 23](#), service router A is connected to the ATM switch or 7670 RSP through an OCx ATM 1 link. This link is configured as the working circuit. Service router B is connected to the same ATM switch or 7670 RSP through an OCx ATM 2 link. This link is configured as the protection circuit.

Figure 23 Multi-Chassis APS Application

IG0001

Communication between service routers A and B is established through link 3. This link is for signaling. To guarantee optimum fail-over time between service routers A and B, link 3 must be a direct physical link between routers A and B.

Sample APS Application: MC-APS with VLL Redundancy

Support of MC-APS to ATM VLLs and Ethernet VLL with ATM SAPs allows MC-APS to operate with pseudowire redundancy in a similar manner that MC-LAG operates with pseudowire redundancy.

The combination of these features provides a solution for access node redundancy and network redundancy as shown in [Figure 24](#).

MC-APS groups are configured as follows:

- MC-APS group between the MSAN on the left and Aggregation Nodes A & B
- MC-APS group between the MSAN on the right and Aggregation Nodes C & D

Figure 24 Access and Node and Network Resilience





Fig_3

An example of a customer application in the mobile market is shown in [Figure 25](#).

Figure 25 MC-APS with ATM VLL Redundancy


OSSG145

In the application shown in [Figure 25](#), 2G and 3G cell sites are aggregated into a Tier 2 or Tier 3 hub site before being backhauled to a Tier 1 site where the radio network controller (RNC) which terminates user calls is located. This application combines MC-APS on the RNC access side and pseudowire redundancy and pseudowire switching on the core network side. pseudowire switching is used in order to separate the routing domains between the access network and the core network.

Sample APS Application: RAN Aggregation with Microwave Radio Transport

[Figure 26](#) shows a RAN aggregation network deployment example. In this example Uni-dir 1+1 Sig+Data APS is being used.

Figure 26 Mobile RAN with Microwave Transport Example



OSSG327

As shown in [Figure 26](#), some APS-protected interfaces may require microwave radio transport. [Figure 27](#) shows APS-protected links between two routers that use Microwave transport. The radio equipment acts as a SONET section/SDH regenerator section equipment, yet it implements Unidirectional APS-like processing to provide equipment protection on the local/remote radio sites respectively.

The active RX line signal (switched independently from TX) is being transmitted over the radio link to the far end radio where the signal gets transmitted on both active and inactive circuits.

The radio reacts on APS triggered failures as detected by the segment termination function: LOS, LOF, manual APS commands, and optionally BER SF/SD. Since the radio does not terminate the SONET/SDH line layer, any line signaling (including kbytes signaling for APS, line alarms like RDI/AIS) are not terminated by the radio and arrive at a far-end router.


Note that the far-end router can either send line alarms based on its active link status or based on physical circuit status (for example, an L-RDI with valid data will be received on the router).

To facilitate a deployment, as shown in this example, some of following features of the 7750 SR-c12 routers are employed:

- Uni-dir 1+1 Sig+Data APS switching mode.
- Configurable L-RDI suppression.

- Active RX circuits are selected based on local conditions only. The SONET K Bytes are not needed to coordinate switch actions, but they are still used since they flow through and reach the far-end router.
- Ports are not failed on L-RDI, as L-RDI may be received on both ports momentarily, as a result of a local radio APS switch or, permanently as a result of a remote router APS switch (with remote radio selecting traffic from the TX line on the same port as failed RX line on the router).
- For some radio equipment, a radio can cause an APS switch resulting in the far end radio detecting radio alarm and generating L-AIS toward its locally attached router on both circuits. In some cases, that router also detects BER SD/BER SF conditions on both circuits as well. Therefore, to localize failure recovery, the 7750 SR-c12 can optionally debounce those alarms so a remote router does not invoke an APS switch on a local failure condition.

Figure 27 1+1 APS Protected Microwave SDH Transport



2.3.2.14 Inverse Multiplexing Over ATM (IMA)

IMA is a cell based protocol where an ATM cell stream is inverse-multiplexed and de-multiplexed in a cyclical fashion among ATM-supporting channels to form a higher bandwidth logical link where the logical link concept is referred as an IMA group. By grouping channels into an IMA group, customers gain bandwidth management capability at in-between rates (for example, between E-1/DS-1 and E-3/DS-3 respectively) through addition/removal of channels to/from the IMA group.

In the ingress direction, traffic coming over multiple ATM channels configured as part of a single IMA group, is converted into a single ATM stream and passed for further processing to the ATM Layer where service-related functions, for example Layer 2 TM, or feeding into a pseudowire are applied. In the egress direction, a single ATM stream (after service functions are applied) is distributed over all paths that are part of an IMA group after ATM layer processing takes place.

An IMA group interface compensates for differential delay and allows only for a minimal cell delay variation. The interface deals with links that are added, deleted or that fail. The higher layers see only an IMA group and not individual links, therefore service configuration and management is done using IMA groups, and not individual links that are part of it.

The IMA protocol uses an IMA frame as the unit of control. An IMA frame consists of a series of consecutive (128) cells. In addition to ATM cells received from the ATM layer, the IMA frame contains IMA OAM cells. Two types of cells are defined: IMA Control Protocol (ICP) cells and IMA filler cells. ICP cells carry information used by IMA protocol at both ends of an IMA group (for example IMA frame sequence number, link stuff indication, status and control indication, IMA ID, TX and RX test patterns, version of the IMA protocol, and so on). A single ICP cell is inserted at the ICP cell offset position (the offset may be different on each link of the group) of each frame. Filler cells are used by the transmitting side to fill up each IMA frame in case there are not enough ATM stream cells from the ATM layer, so a continuous stream of cells is presented to the physical layer. Those cells are then discarded by the receiving end. IMA frames are transmitted simultaneously on all paths of an IMA group and when they are received out of sync at the other end of the IMA group link, the receiver compensates for differential link delays among all paths.

2.3.2.14.1 Inverse Multiplexing over ATM (IMA) Features

Hardware Applicability

IMA is supported on channelized ASAP MDAs.

Software Capabilities

Nokia's implementation supports IMA functionality as specified in ATM Forum's Inverse Multiplexing for ATM (IMA) Specification Version 1.1 (af-phy-0086.001, March 1999). The following capabilities are supported:

- TX Frame length — Only IMA specification default of 128 cells is supported.

-
- IMA version — Both versions 1.0 and 1.1 of IMA are supported. There is no support for automatically falling to version 1.0 if the far end advertises 1.0 support, and the local end is configured as 1.1. Due to potential protocol interoperability issues between IMA 1.0 implementations, it is recommended that IMA version 1.1 is used whenever possible.
 - Alpha, beta, and gamma values supported are defaults required by the IMA specification (values of 2, 2, and 1 respectively).
 - Clock mode — Only IMA specification default of common clock mode is supported (CTC).
 - Timing reference link — The transmit timing reference link is chosen first among the active links in an IMA group. If none found, then it is chosen among the usable links or finally, among the unusable links.
 - Cell Offset Configuration — The cell offsets for IMA links are not user configurable but internally assigned according to the recommended distribution described in the IMA spec.
 - TX IMA ID — An internally assigned number equal to the IMA bundle number.
 - Minimum Links — A configurable value is supported to control minimum member links required to be up for an IMA group to stay operationally up.
 - Maximum Group Bandwidth — A configurable value is supported to specify maximum bandwidth available to services over an IMA group. The maximum may exceed the number of minimum/configured/active links allowing for overbooking of ATM shaped traffic.
 - Symmetry mode — Only IMA specification default of symmetric operation and configuration is supported.
 - Re-alignment — Errors that require a re-alignment of the link (missing or extra cells, corrupted frame sequence numbers), are dealt with by automatically resetting the IMA link upon detection of an error.
 - Activation/Deactivation Link Delay Timers — Separate, configurable timers are supported defining the amount of delay between detection of LIF, LODS and RFI-IMA change and raising/clearing of a respective alarm to higher layers and reporting RXIFailed to the far end. This protocol dampening mechanism protects those higher layers from bouncing links.
 - Differential delay — A configurable value of differential delay that will be tolerated among the members of the IMA group is supported. If a link exceeds the configured delay value, then LODS defect is declared and protocol management actions are initiated as required by the IMA protocol and as governed by Link Activation and Deactivation procedures. The differential delay of a link is calculated based on the difference between the frame sequence number received on the link and the frame sequence number received on the fastest link (a link on which the IMA frame was received first).

- Graceful link deletion — The option is supported for remotely originated requests only. To prevent data loss on services configured over an IMA group, it is recommended to initiate graceful deletion from the far end before a member link is deleted or a physical link is shutdown.
- IMA test pattern — Nokia's implementation supports test pattern procedures specified in the IMA specification. Test pattern procedures allow debugging of IMA group problems without affecting user data. Test pattern configurations are not preserved upon a router reboot.
- Statistics — Nokia's IMA implementation supports all standard-defined IMA group and IMA link status and statistics through proprietary TIMETRA-PORT-MIB. Display and monitoring of traffic related interface/SAP statistics is also available for IMA groups and services over IMA groups on par with physical ATM interfaces and services.
- Scaling — Up to 8 member links per IMA group, up to 128 groups per MDA and all DS-1/E-1 links configurable per MDA in all IMA groups per MDA are supported.

2.3.2.15 Ethernet Local Management Interface (E-LMI)

The Ethernet Local Management Interface (E-LMI) protocol is defined in Metro Ethernet Forum (MEF) technical specification MEF16. This specification largely based on Frame Relay - LMI defines the protocol and procedures that convey the information for auto-configuration of a CE device and provides the means for EVC status notification. MEF16 does not include link management functions like Frame Relay LMI does. In the Ethernet context that role is already accomplished with Clause 57 Ethernet OAM (formerly 802.3ah).

The SR OS currently implements the User Network Interface-Network (UNI-N) functions for status notification supported on Ethernet access ports with dot1q encapsulation type. Notification related to status change of the EVC and CE-VLAN ID to EVC mapping information is provided as a one to one between SAP and EVC.

The E-LMI frame encapsulation is based on IEEE 802.3 untagged MAC frame format using an ether-type of 0x88EE. The destination MAC address of the packet 01-80-C2-00-00-07 will be dropped by any 802.1d compliant bridge that does not support or have the E-LMI protocol enabled. This means the protocol cannot be tunneled.

Status information is sent from the UNI-N to the UNI-C, either because a status inquiry was received from the UNI-C or unsolicited. The Active and Not Active EVC status are supported. The Partially Active state is left for further study.

The bandwidth profile sub-information element associated with the EVC Status IE does not use information from the SAP QoS policy. A value of 0 is used in this release as MEF 16 indicates the bandwidth profile sub-IE is mandatory in the EVC Status IE. The EVC identifier is set to the description of the SAP and the UNI identifier is set to the description configured on the port. Further, the implementation associates each SAP with an EVC. Currently, support exists for CE-VLAN ID/EVC bundling mode.

The E-LMI the UNI-N can participate in the OAM fault propagation functions. This is a unidirectional update from the UNI-N to the UNI-C and interacting with service manager of VLL, VPLS, VPRN and IES services.

2.3.2.16 Link Layer Discovery Protocol (LLDP)


The IEEE 802.1ab Link Layer Discovery Protocol (LLDP) standard defines protocol and management elements that are suitable for advertising information to stations attached to the same IEEE 802 LAN (emulation) for the purpose of populating physical or logical topology and device discovery management information databases. The protocol facilitates the identification of stations connected by IEEE 802 LANs/MANs, their points of interconnection, and access points for management protocols.

Note that LAN emulation and logical topology wording is applicable to customer bridge scenarios (enterprise/carrier of carrier) connected to a provider network offering a transparent LAN emulation service to their customers. It helps the customer bridges detect misconnection by an intermediate provider by offering a view of the customer topology where the provider service is represented as a LAN interconnecting these customer bridges.

The IEEE 802.1ab standard defines a protocol that:

- Advertises connectivity and management information about the local station to adjacent stations on the same IEEE 802 LAN.
- Receives network management information from adjacent stations on the same IEEE 802 LAN.
- Operates with all IEEE 802 access protocols and network media.
- Establishes a network management information schema and object definitions that are suitable for storing connection information about adjacent stations.
- Provides compatibility with a number of MIBs as shown in [Figure 28](#).


Figure 28 LLDP Internal Architecture for a Network Node



OSSG262

Network operators must be able to discover the topology information in order to detect and address network problems and inconsistencies in the configuration. Moreover, standard-based tools can address the complex network scenarios where multiple devices from different vendors are interconnected using Ethernet interfaces.

The example shown in [Figure 29](#) depicts a MPLS network that uses Ethernet interfaces in the core or as an access/handoff interfaces to connect to different kind of Ethernet enabled devices such as service gateway/routers, QinQ switches, DSLAMs or customer equipment.

Figure 29 Generic Customer Use Case For LLDP

IEEE 802.1ab LLDP running on each Ethernet interfaces in between all the above network elements may be used to discover the topology information.

Operators who are utilizing IOM/IMM can tunnel the nearest-bridge at the port level using the **tunnel-nearest-bridge** command under the **config>port>ethernet>lldp>destmac** (nearest-bridge) hierarchy. The dest-mac nearest-bridge must be disabled for tunneling to occur.

2.3.2.16.1 LLDP Protocol Features

LLDP is a unidirectional protocol that uses the MAC layer to transmit specific information related to the capabilities and status of the local device. Separately from the transmit direction, the LLDP agent can also receive the same kind of information for a remote device which is stored in the related MIBs.

LLDP itself does not contain a mechanism for soliciting specific information from other LLDP agents, nor does it provide a specific means of confirming the receipt of information. LLDP allows the transmitter and the receiver to be separately enabled, making it possible to configure an implementation so the local LLDP agent can either transmit only or receive only, or can transmit and receive LLDP information.

The information fields in each LLDP frame are contained in an LLDP Data Unit (LLDPDU) as a sequence of variable length information elements, that each include type, length, and value fields (known as TLVs), where:

- Type identifies what kind of information is being sent.
- Length indicates the length of the information string in octets.
- Value is the actual information that needs to be sent (for example, a binary bit map or an alphanumeric string that can contain one or more fields).

Each LLDPDU contains four mandatory TLVs and can contain optional TLVs as selected by network management:

- Chassis ID TLV
- Port ID TLV
- Time To Live TLV
- Zero or more optional TLVs, as allowed by the maximum size of the LLDPDU
- End Of LLDPDU TLV

The chassis ID and the port ID values are concatenated to form a logical identifier that is used by the recipient to identify the sending LLDP agent/port. Both the chassis ID and port ID values can be defined in a number of convenient forms. Once selected however, the chassis ID/port ID value combination remains the same as long as the particular port remains operable.

A non-zero value in the TTL field of the Time To Live TLV tells the receiving LLDP agent how long all information pertaining to this LLDPDU's identifier will be valid so that all the associated information can later be automatically discarded by the receiving LLDP agent if the sender fails to update it in a timely manner. A zero value indicates that any information pertaining to this LLDPDU's identifier is to be discarded immediately.

A TTL value of zero can be used, for example, to signal that the sending port has initiated a port shutdown procedure. The End Of LLDPDU TLV marks the end of the LLDPDU.

The implementation defaults to setting the port-id field in the LLDP OAMPDU to **tx-local**. This encodes the port-id field as ifIndex (sub-type 7) of the associated port. This is required to support some releases of SAM. SAM may use the ifIndex value to properly build the Layer Two Topology Network Map. However, this numerical value is difficult to interpret or readily identify the LLDP peer when reading the CLI or MIB value without SAM. Including the **port-desc** option as part of the **tx-tlv** configuration allows a Nokia remote peer supporting **port-desc** preferred display logic (11.0r1) to display the value in the port description TLV instead of the port-id field value. This does not change the encoding of the port-id field. That value continues to represent the ifIndex. In some environments, it may be important to select the specific port information that is carried in the port-id field. The operator has the ability to control the encoding of the port-id information and the associated subtype using the **port-id-subtype** option. Three options are supported for the port-id-subtype:

tx-if-alias — Transmits the ifAlias String (subtype 1) that describes the port as stored in the IF-MIB, either user configured description or the default entry (i.e.10/100/Gig Ethernet SFP)

tx-if-name — Transmits the ifName string (subtype 5) that describes the port as stored in the IF-MIB, ifName info.

tx-local — The interface ifIndex value (subtype 7)

IPv6 (address subtype 2) and IPv4 (address subtype 1) LLDP System Management addresses are supported. The IP addresses can be selected from the system IP addressing, the out-of-band management address (BOF), or both.

2.3.2.17 Exponential Port Dampening

Exponential Port Dampening (EPD) provides the ability to automatically block a port from re-use for a period of time after physical link-down and physical link-up events. If a series of link-down and -up events occur close together, EPD keeps the port's operational state down for a longer period than if only one down-up event occurred. This keeps the router from using that port if there are external events causing the link state to fluctuate. The more events that occur, the longer the port is kept down and avoided by the routing protocols.

EPD behavior uses a fixed penalty amount per link-down event and a half-life decay equation to reduce these penalties over time. Exponential decay is defined by the following equation:

$$N(t) = N_0 \left(\frac{1}{2} \right)^{\frac{t}{t_{1/2}}}$$

sw0109

where:

$N(t)$ is the quantity that still remains after a time t

N_0 is the initial quantity


$t_{1/2}$ is the half-life

In dampening, N_0 refers to the starting penalties from the last link-down event. The quantity $N(t)$ refers to the decayed penalties at a specific time, and is calculated starting from the last link-down event (that is, from the time when N_0 last changed).

This equation can also be used on a periodic basis by updating the initial quantity value N_0 each period and then computing the new penalty over the period (t).

An example of the use of this feature is shown in [Figure 30](#).

Figure 30 Exponential Port Dampening Example



At time ($t = 0$), the initial condition has the link up, the accumulated penalties are zero, the dampening state is idle, and the port operational state is up. The following series of events and actions occur.

1. $t = 5$: Link-down event
 - i. the accumulated penalties are incremented by 1000
 - ii. the accumulated penalties now equal 1000, which is less than the suppress threshold (of 1500), so the dampening state is idle
 - iii. because the dampening state is idle, link-down is passed to the upper layer
 - iv. link-down triggers the port operational state to down
2. $t = 9$: Link-up event
 - i. the accumulated penalties equal 869, which is less than the suppress threshold, so the dampening state remains as idle
 - ii. because the dampening state is idle, link-up is passed to the upper layer
 - iii. link-up triggers the port operational state to up
3. $t = 13$: Link-down event
 - i. the accumulated penalties are incremented by 1000
 - ii. the accumulated penalties now equal 1755, which is greater than the suppress threshold, so the dampening state is changed to active
 - iii. because the dampening state just transitioned to active, link-down is passed to the upper layer
 - iv. link-down triggers the port operational state to down
4. $t = 17$: Link-up event
 - i. the accumulated penalties equal 1527, which is above the reuse threshold (of 1000) and greater than the suppress threshold, so the dampening state remains as active
 - ii. because the dampening state is active, link-up is not passed to the upper layer
 - iii. the port operational state remains down
5. $t = 21$: Link-down event
 - i. the accumulated penalties are incremented by 1000
 - ii. the accumulated penalties now equal 2327, which is above the reuse threshold, so the dampening state remains as active
 - iii. because the dampening state is active, link-down is not passed to the upper layer
 - iv. the port operational state remains down
6. $t = 25$: Link-up event

- i. the accumulated penalties equal 2024, which is above the reuse threshold, so dampening state remains as active
 - ii. because the dampening state is active, link-up is not passed to the upper layer
 - iii. the port operational state remains down
- 7. $t = 46$: Accumulated penalties drop below the reuse threshold
 - i. the accumulated penalties drop below the reuse threshold, so the dampening state changes to idle
 - ii. because the dampening state is idle and the current link state is up, link-up is passed to the upper layer
 - iii. the port operational state changes to up
- 8. $t = 94$ to 133 : Link-down and link-up events every second
 - i. similar to previous events, the accumulated penalties increment on every link-down event
 - ii. the dampening state transitions to active at $t = 96$, and link state events are not sent to the upper layer after that time
 - iii. the upper layer keeps the port operational state down after $t = 96$
 - iv. the accumulated penalties increment to a maximum of 4000
- 9. $t = 133$: Final link event of link-up
 - i. the accumulated penalties equal 3863
 - ii. the dampening state remains active and link state events are not sent to the upper layer
 - iii. the upper layer keeps the port operational state down
- 10. $t = 172$: Accumulated penalties drop below the reuse threshold
 - i. the accumulated penalties drop below the reuse threshold, so the dampening state changes to idle
 - ii. because the dampening state is idle and the current link state is up, link-up is passed to the upper layer
 - iii. the port operational state changes to up

2.3.3 Per Port Aggregate Egress Queue Statistics Monitoring

Monitoring the aggregate egress queue statistics per port provides in-profile, out-of-profile, and total statistics for both forwarded and dropped packets and octets on a given port.

When enabled, all queues on the port are monitored, including SAP egress, network egress, subscriber egress, and egress queue group queues, as well as system queues which can be used, for example, to send port-related protocol packets (LACP, EFM, and so on).

This is enabled and disabled using the following command:

```
config port <port-id> [no] monitor-agg-egress-queue-stats
```

When enabled, the line card will poll the related queues to derive the aggregates which provide the delta of the queue statistics since turning on the monitoring. This means that the reported statistics are not reduced by those from a deleted queue and so the aggregates correctly represent the forwarded/dropped statistics since the start of monitoring.

The aggregates can be shown with the following command:

```
show port [<port-id>] [statistics [egress-aggregate]] [detail]
```

As an example, the output below enables monitoring of aggregate egress queue statistics on port 2/1/1 and then shows the monitored statistics:

```
*A:PE# configure port 2/1/1 monitor-agg-egress-queue-stats
*A:PE# show port 2/1/1 statistics egress-aggregate
```

```
=====
Port 2/1/1 Egress Aggregate Statistics on Slot 2
=====
```

	Forwarded	Dropped	Total
PacketsIn	144	0	144
PacketsOut	0	0	0
OctetsIn	12353	0	12353
OctetsOut	0	0	0

```
=====
*A:PE#
```

To clear the aggregate statistics, the monitoring must be disabled and then re-enabled. The aggregate statistics are also cleared when the card is cleared (using a **clear card slot-number** command) or power-cycled (with the **tools perform card slot-id** command). Additionally, aggregate statistics related to MDA are cleared when the MDA is cleared (using the **clear mda mda-id** command) or the MDA is inserted into an IOM. The aggregate statistics are not cleared when a **shutdown/no shutdown** is performed on the card and/or MDA.

There is no specific limit on the number of queues that can be monitored, but the amount of each line card's CPU resources allocated to the monitoring is bounded; consequently, when more queues on a card's ports are monitored, the aggregate statistics will be updated the less frequently.

Monitoring of aggregate statistics is supported on **PXC** sub-ports but not on a **PXC** physical port. It is also not supported on satellite ports or ports on an HSMDA.

2.4 FP4 Data Path Mapping

The mapping between a card and its MDAs, FPs, MACs, connectors, and ports on FP4-based hardware can be displayed using the **show datapath** command.

```
A:PE# show datapath 1/1 mac 2 detail
```


```
=====
Card      MDA      FP      MAC Chip Num  Connector  Port
-----
1          1        1        2             c7         1/1/c7/1
1          1        1        2             c8         1/1/c8/1
1          1        1        2             c9         1/1/c9/1
1          1        1        2             c10        1/1/c10/1
1          1        1        2             c11        1/1/c11/1
1          1        1        2             c12        1/1/c12/1
=====
*A:PE#
```

2.5 Port Cross-Connect (PXC)

Port Cross-Connect (PXC) functionality refers to a method of redirecting ingress traffic received on I/O ports to an operator-chosen forwarding complex (anchor point). This traffic is then pre-processed in the egress data path of the anchor card and consequently looped back into the ingress data path on the same anchor card. This is shown in [Figure 31](#). Pre-processing of the traffic is necessary in some cases due to its complex nature. Examples of traffic pre-processing include removal of an encapsulation layer in the packet or modifying IP header fields that cannot be performed in the ingress data path. With this, the service termination point is moved from the I/O port (where traffic is received on a node) to a PXC port on the anchor card. Looping the traffic from egress to ingress on the anchor point is performed by a physical port in a loopback mode. This port is referred to as a PXC port or simply as a PXC.

Conceptually, PXC functionality is similar to the functionality provided by two externally interconnected physical ports where the packets leave the system through one port and then immediately looped back into another port through an external cable. [Figure 31](#) depicts the steps involved in traffic preprocessing that utilizes a PXC:

- Traffic entering a node through an I/O port is processed by the local ingress forwarding path (1) on the line cards 1 and 2. Traffic is then directed (1) toward the PXC port (3) on the line card 3.
- Additional pre-processing is performed in the egress forwarding path (2) in line card 3 just before traffic is sent to the PXC port 3.
- The PXC loops the traffic toward the local ingress forwarding path (4) where it is further processed.

Figure 31 Traffic Pre-Processing Utilizing Cross-Connected Port (PXC)

0941

2.5.1 PXC Terminology

Port Cross-Connect (PXC) or pxc port — This is a physical port that is internally looped to connect the egress forwarding path to the ingress forwarding path within the same line card.


PXC sub-port — This is a logical port that is created under the PXC port. Two PXC sub-ports are created per PXC and represent an upstream path and a downstream path. This is shown in [Figure 31](#).

Anchor Point — This refers to a location in the node where a PXC or a group of PXCs are created. An anchor point can represent a PXC on a single line card or a group of line cards where PXCs are distributed over multiple line cards in a LAG fashion for redundancy purposes.

2.5.2 PXC - Physical Port in Cross-Connect (Loopback) Mode

The concept of a PXC port on an FP3-based line card is shown in [Figure 32](#). The PXC does not require an optical transceiver.

Figure 32 PXC



The physical port is placed in a cross-connect mode with the following commands:

```
configure
  port-xc
    pxc <pxc-id> create
    description <string>
    port <port-id>
    [no] shutdown
```

Once a physical port is associated with a PXC ID using the above commands, the node will automatically create a corresponding pair of mated PXC sub-ports (these are logical ports representing upstream/downstream paths under the PXC). The PXC is auto-configured as a hybrid port with MTU preset to 9212 bytes, and the encapsulation set to dot1q and dot1x tunneling is turned on.

These parameters cannot be changed once the physical port becomes cross-connected.

The following applies to PXC ports:

- Only unused physical ports (not associated with an interface or SAP) can be referenced inside of a PXC ID configuration.
- Once inside the PXC ID configuration, the physical port cannot be removed from that PXC ID if the corresponding PXC sub-ports are currently in use.

- Once inside the PXC ID configuration, the physical port cannot be used outside of the PXC context. For example, a regular IP interface cannot use this physical port, or a SAP on that port cannot be associated with a service.
- A physical port can be associated with only one PXC ID configuration.

2.5.2.1 Operational State

The operational state of the PXC is dependent on its administrative state. When a PXC is operationally up, the port status LED on the faceplate blinks amber. The port activity LED will be lit green in the presence of traffic on PXC ports and will turn off in absence of traffic on PXC ports.

The presence of the optical transceiver on the PXC has no effect on its operational state. Traffic cannot be sent out through the transceiver or be received from the transceiver from the outside. However, the existing traps related to insertion or removal of a transceiver (SFF Inserted/Removed), are supported.


The “Signal-Fail” alarm on the PXC is suppressed.

The operational state of the PXC ID is derived from its administrative state which is the operational state of the underlying physical port and the admin state of the corresponding PXC sub-ports.

2.5.3 PXC Sub-Ports

To explain the need for PXC sub-ports under an internally cross-connected physical port, an analogy with two distinct physical ports that are connected via external cable is shown in [Figure 33](#).

Figure 33 Two Cross-Connected Ports vs Single Cross-Connected Port



0943

Bidirectional connectivity provided by the PXC requires two sub-ports (one in each direction). These sub-ports are referred to as PXC sub-ports and are used by the router as logical configurations to transmit traffic in both directions over a single physical port that is internally cross-connected.

Once the physical port is associated with a PXC ID, a pair of mated PXC sub-ports is automatically created by the system:

```
configure
  port pxc-<pxc-id>.a
  shutdown
```

```
port pxc-<pxc-id>.b
shutdown
```

The two PXC sub-ports are distinguishable by .a and .b suffixes. They transmit traffic toward each other, thus simulating two ports that are interconnected.

The PXC sub-ports can be accessed through the configuration in order to modify their parameters (QoS, and so on). However, certain PXC parameters are fixed and cannot be changed. For example, PXC sub-ports are created in a hybrid mode and therefore cannot be modified.

Each PXC sub-port is internally (within the system) represented by an internal four byte VLAN tag. Traffic carried over the PXC will contain four extra bytes which is accounted for in QoS provisioning.

Configuration example:

```
configure
port-xc
  pxc 1 create
    port 1/1/1
    [no] shutdown
  pxc 2 create
    port 1/1/2
    [no] shutdown
```

The configuration shown above automatically creates the following:

```
configure
port
  pxc-1.a      → cross-connected with pxc-1.b
  pxc-1.b      → cross-connected with pxc-1.a
  pxc-2.a      → cross-connected with pxc-2.b
  pxc-2.b      → cross-connected with pxc-2.a
```

2.5.3.1 PXC Sub-Port Operational State

At creation time, the administrative state of the PXC sub-ports is set to shutdown.

The operational state of the PXC sub-ports is dependent on the operational state of the underlying physical port and the state of the corresponding PXC ID.

2.5.4 Port Statistics

There are two types of statistics that can be collected on a regular (non PXC) Ethernet port:

- **Physical port level statistics** provide information about conditions on the data-link layer and physical level, for example, the aggregate number of forwarded and dropped octets/bytes on a MAC level, FCS errors, number of collisions, and so on. These statistics can be viewed with the **show port *port-id*** command and they can be collected into a local file in XML format. Local file collection is enabled with either of the following commands:
 - **config>port>ethernet>collect-stats**
 - **config>port>ethernet>accounting-policy *policy-id***
- **Network level statistics** provide information about forwarded and dropped octets and packets on a per queue level on network ports. These statistics can be viewed with the **show port *port-id* detail** command and they can be collected into a local file in XML format. Local file collection is enabled with either of the following commands:
 - **config>port>ethernet>network>collect-stats**
 - **config>port>ethernet>network>accounting-policy *policy-id***

PXC ports will inherit some of those statistics.

2.5.4.1 Statistics on Physical PXC Ports

The statistics on physical PXC ports are maintained only on MAC level. The internal q-tag used for PXC sub-port identification within the router is included in the displayed octet count. The port-level statistics on a physical PXC port represent the combined upstream and downstream traffic carried by the corresponding PXC sub-ports.

For example, in port level statistics output for a physical PXC port, the **output** count represents the upstream and downstream traffic flowing 'out' of the physical port while the **input** count represents the same looped traffic 'returning into' the same port.

```
=====
Traffic Statistics
=====
                                Input                Output
-----
Octets                        290164703            290164703
Packets                       2712661             2712661
Errors                          0                      0
```

Statistics are cleared when a physical port enters or leaves the PXC mode.

Statistics collection to a local file is not supported on physical PXC ports.


Queues are not instantiated on physical PXC ports and therefore, the network level (queue) statistics are not maintained on physical PXC ports.

2.5.5 LAG with PXC Ports – PXC LAG

PXC sub-ports can be aggregated to form a PXC LAG for increased capacity and anchor point redundancy.

A logical concept of a PXC LAG is shown in [Figure 34](#).

Figure 34 Logical Concept of a LAG on PXC Ports



The LAGs on physical PXC ports must be configured in pairs. A pair of PXC LAGs (LAG1 and LAG2) as shown in this example:

```
configure
  lag 1
    description "lag in the up direction"
    port pxc-1.a
    port pxc-2.a

  lag 2
    description "lag in the down direction"
    port pxc-1.b
    port pxc-2.b
```

Within the router, the two sides of the PXC LAG (LAG 1 and LAG 2 in above example) are not aware that they are cross-connected. As a result, the operational state of one side of the PXC LAG will not be influenced by the state of the PXC LAG on the other side.


PXC sub-ports in a LAG must have the same properties (such as the same speed, and so on). Mixing PXC sub-ports and non-PXC ports is not allowed. The first port added to a LAG will determine the LAG port-type (PXC or non-PXC).

Statistics in the output of the **show lag id statistics** command represent combined traffic carried over the referenced **lag id** and its pair (LAG 1 and LAG 2 in the above example).

2.5.6 Basic PXC Provisioning

The CLI flow is shown in [Figure 35](#). The oval marked 'Operator' represents a configuration step that must be performed by the operator. The block marked 'Dynamic' represents a step that is performed automatically by the system without an operator's assistance.

Figure 35 CLI Flow



0945

2.5.7 QoS

Upstream and downstream traffic on a PXC is funneled through the same physical port. In other words, upstream and downstream traffic is transmitted and then received by the same port. Each forwarding direction (upstream and downstream) is represented by a corresponding PXC sub-port (.a or .b). This is fundamental difference from a non-PXC port which transmits traffic only in one direction.

Traffic traversing a PXC contains an overhead of four bytes per packet that are attributed to the internal VLAN tag used for PXC sub-port identification within the 7750 SR. However, these four bytes are not accounted in configured QoS rates. Therefore, the operator should take this into consideration when configuring rates on QoS objects under PXC ports.

Funneling traffic from two PXC sub-ports through a single physical PXC ports requires separate bandwidth management per PXC sub-port. The sum of the configured bandwidth maximums for the egress port scheduler (EPS) under the two PXC sub-ports should not exceed the bandwidth capacity of the underlying physical port minus the overhead introduced by the four bytes attributed to the internal VLAN tag.

If the sum of the configured EPS bandwidth under the two PXC sub-ports exceeds the bandwidth capacity of the underlying physical port, the traffic scheduling, if there is PXC congestion, will be performed according to the configured queue types (expedited, non-expedited) and their configured CIR/PIR rates. However, then the configured EPS bandwidth limits will not be honored.

2.5.7.1 Queue Allocation on PXC Sub-Ports

PXC sub-ports are auto-configured in a hybrid mode and cannot be changed via configuration. They each have a set of queues on the network egress and a set of queues on the access egress/ingress (per SAP or ESM subscriber). Queues on the network ingress are shared per MDA, just as they are on non-PXC ports in the hybrid mode.

Queue groups are allocated per PXC sub-ports.

2.5.7.2 Pool Allocations on PXC Ports

Queue buffers are created in buffer pools and are used for traffic buffering when queues are congested. Buffer pools are allocated per forwarding complex or per PXC physical port.

Each physical PXC port has three associated buffer pools:

- access ingress
- access egress
- network egress

The network ingress pool is shared between all physical ports on a forwarding complex. The size of the buffer pools is automatically determined by the system based on the forwarding complex type and port configuration.

2.5.7.3 QoS Summary

The QoS hierarchy supported on a PXC is similar to the existing QoS hierarchy supported on regular (non-PXC) Ethernet ports.

A notable difference between PXC and non-PXC ports in the area of QoS is the bandwidth control per physical port. In the PXC. The sum of the configured bandwidth maximums for the two Ethernet port schedulers applied to the PXC sub-ports must not exceed the bandwidth capacity of the underlying physical PXC port. This means that the bandwidth on a PXC should be partitioned per direction (upstream and downstream).

With regular (non-PXC) Ethernet ports, this sum is not relevant in the context of a single physical port. Each non PXC Ethernet port can transmit traffic only in one direction (upstream or downstream), and thus, one EPS is sufficient to control port congestion. The configured bandwidth of the EPS should not exceed the bandwidth capacity of the physical port.

2.5.8 Mirroring and LI on PXC Ports

Traffic on a PXC sub-port can be mirrored or lawfully intercepted (LI). For example, subscriber “Annex1” traffic arriving on a PXC sub-port will be mirrored if “Annex1” is configured as a mirror or LI source. A PXC sub-port can also be used to transmit mirror/LI traffic out from a mirror-destination service (such as a mirror-dest SAP or SDP can egress out a PXC sub-port, or a routable LI encapsulated packet can be forwarded and transmitted out a PXC sub-port).

A mirror destination can be configured to transmit mirrored and LI traffic out a SAP on a PXC sub-port that is then cross connected into a VPLS service where a VXLAN encapsulation is then added to the mirrored packets before transmission out of the node.

The internal q-tag that represent the PXC sub-port within the system (a satellite or anchor tag) will be included in the lawfully intercepted copy of the packet for traffic intercepted (mirrored) on the ingress side of a PXC sub-port when the associate mirror-dest service is of type **ether** (the default) with routable lawful interception the encapsulation (**mirror-dest>encap**).

2.5.9 Multi-Chassis Redundancy

Multi-Chassis Synchronization (MCS) configuration (**config>redundancy>multi-chassis>peer>sync**) is supported for entities utilizing PXC's. However, MC-LAG is not supported directly on PXC's since PXC are not directly connected to external equipment. MC-LAG is supported on I/O ports that are front ending PXC ports.

2.5.10 Health Monitoring on the PXC Sub-Ports

Health monitoring on the PXC ports is based on the following:

- CRC monitoring (link quality) on the RX side of the port (**config>port>ethernet>crc-monitor**) and/or
- CRC monitoring (link quality) on the path from the IOM toward the MDA (**config>port>ethernet>down-on-internal-error**). Note that the **tx-disable** flag (disable remote laser on error) is not supported on PXC ports since PXC ports are looped.

Health monitoring of the PXC **sub**-ports is based on:

- efm-oam — The Information OAMPDUs are transmitted by each peer (PXC sub-port) at the configured intervals. Their purpose is to perform keepalive and critical notification functions.
- CRC errors are only recorded, if frames are transmitted.
- Crossing the signal degrade (SD) threshold: raise log event

CRC monitoring on the RX side of the PXC ports has the following characteristics:

- Monitors ingress error conditions
- Compares error counts against configurable thresholds
- CRC errors are only recorded if frames are transmitted
- Crossing the signal degrade (SD) threshold raises an log event
- Crossing the signal failure (SF) threshold takes the port operational state down
- Error rate thresholds uses format $m \cdot 10^{-n}$

Both threshold (n) and multiplier (m) are configurable.

Health monitoring on PXC ports is disabled by default.

In addition to the explicitly configured aforementioned health monitoring mechanisms, PXC operational state transitions will be, by default, reported by a port UP/DOWN trap:

```
478 2015/10/22 14:08:15.86 UTC WARNING: SNMP #2004 Base pxc-1.b Interface pxc-1.b is not operational
```

```
478 2015/10/22 14:08:15.86 UTC WARNING: SNMP #2004 Base pxc-1.b Interface pxc-1.b is operational
```

2.5.11 Configuration Example

In the following example, subscriber (ESM) traffic with QinQ encapsulation arriving on two different I/O line cards (3 and 4) is terminated on the PXC LAG on line cards 1 and 2. With this method, if one of the I/O line cards fails, the subscriber traffic remains unaffected (continues to be terminated on line cards 1 and 2) provided that the proper protection mechanism is implemented in the access part of the network. This protection mechanism in the access part of the network must ensure that traffic arriving on card 3 can be rerouted to card 4 if card 3 fails. The opposite must be true as well (path to card 4 must be protected by a path to card 3).

A PXC can be on any card, independent of I/O ports.

The following displays an example of an I/O port configuration on cards 3 and 4:

```
configure
  port 3/1/1
    description "access I/O port on card 3; ecap is null which means that
      all VLAN tagged and untagged traffic will be accepted"
    ethernet
      mode access
      encap-type null
  port 4/1/1
    description "access I/O port on card 4; ecap is null which means that
      all VLAN tagged and untagged traffic will be accepted"
    ethernet
      mode access
      encap-type null
```

The following displays an example of a PXC configuration on cards 1 and 2:

```
configure
  port-xc
    pxc 1 create
      description "PXC on card 1"
      port 1/1/1
      no shutdown
    pxc 2 create
      description "PXC on card 2"
      port 2/1/1
  no shutdown
```

The above configuration segment will trigger automatic creation of a pair of mated PXC sub-ports. This is shown below. The desired sub-port encapsulation must be set manually by the operator (the default is dot1q). PXC sub-ports will transparently pass traffic with preserved QinQ tags from the .b side of the PXC (I/O side) to the .a side of the PXC where *.* capture SAP will be configured.

```
configure
  port pxc-1.a
    description "termination PXC side; *.* capture SAP will be
      configured here"
    encap-type qinq
    no shutdown
  port pxc-1.b
    description "transition PXC side; all VLAN tags (*) will be
      transparently passed via this side"
    encap-type dot1q
    no shutdown
  port pxc-2.a
    description "together with pxc-1.a, this sub-port is a member of
      LAG 1"
    encap-type qinq
    no shutdown
  port pxc-2.b
    description "together with pxc-1.b, this sub-port is a member of
      LAG 2"
    encap-type dot1q
    no shutdown
```


The following displays an example of a PXC LAG configuration:

```
configure
  lag 1 create
    description "terminating side of the cross-connect"
    port pxc-1.a
    port pxc-2.a
  lag 2 create
    description "transient side of the cross-connect"
    port pxc-1.b
    port pxc-2.b
```

Passing traffic from the I/O side on access (ports 3/1/1 and 4/1/1) via the transient PXC sub-ports pxc-1.b and pxc-2.b to the termination side of the PXC is performed via VPLS.

```
configure
  service vpls 1 create customer 1
    description "stitching access side to the anchor"
    split-horizon-group "access (I/O) side" create
    sap 3/1/1 split-horizon-group "access" create
      description "I/O port"
    sap 4/1/1 split-horizon-group "access" create
      description "I/O port"
    sap lag-2:*
      description "transient side od PXC"
```

The following displays an example of capture SAPs on the anchor:

```
configure
  service vpls 3 create customer 1
    description "VPLS with capture SAPs"
    sap lag-1:10.* capture-sap create
      description "termination side of PXC; traffic with
        S-tag=10 will be extracted here"
      trigger-packet dhcp dhcpv6 pppoe

    sap lag-1:11.* capture-sap create
      description "termination side of PXC; traffic with
        S-tag=11 will be extracted here".
```

2.6 Forwarding Path Extensions (FPE)

Certain applications in the SR OS require extra traffic processing in the forwarding plane. Such additional traffic processing is facilitated by an internal cross-connect that utilizes PXC ports (described in the [Port Cross-Connect \(PXC\)](#)). Application-specific use of the cross-connect is built on the common premise that the traffic must be steered from the input ports to the PXC ports where the traffic can be looped for additional processing in the forwarding plane. To shield the operator from the intricacies involved when configuring application-specific cross-connect attributes, a CLI construct referred to as Forwarding Path Extensions (FPE) simplifies provisioning of various applications which rely on PXC functionality.

Two examples of applications which rely on PXC and FPE are:

- Anchored PW-ports where PW payload termination in Layer 3 services is disjointed from I/O ports in the system.
- VXLAN termination on non-system IPv4 addresses and VXLAN IPv6 underlay.


Application-specific uses of PXC ports and FPEs are described in the respective user guides (*7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide, 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN, and 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*).

The FPE configuration provides information to the SR OS node necessary to associate the application with the PXC (paired PXC sub-ports or PXC based LAG ids). Consequently, the SR OS node will set up the internal logic utilizing PXC as required by the application.

An example of FPE provisioning is given in [Figure 36](#).

- The first three steps are applicable to PXC port provisioning.
- Association between the application and the PXC is performed in steps 4 and 5. In this particular example, two applications can be configured: PW-port or VXLAN-termination (non-system IPv4 termination or IPv6 underlay). These applications require internal configuration of SDPs and their IDs are allocated from the user configurable range. In order to prevent conflict between the user provisioned SDP ids and internally configured SDP id in FPE case, a range of SDP ids that will be used by FPE is reserved by the **sdp-id-range** commands under the **config>fwd-path-ext** CLI hierarchy.
- Application-specific configuration is performed in step 6, partially by the operator and partially by the system. This is described in the application-specific user guides.

Figure 36 FPE - Sample Provisioning Steps



Once the PXC sub-port or LAG is associated with an FPE object, the manual creation (by the operator) of IP interfaces and SAPs under such PXC sub-ports/LAGs is not permitted. Only the internal SR OS system is allowed to reference these PXC sub-ports/LAGs in internal IP interfaces and SAPs, as required by each application.

However, the PXC sub-ports and LAG parameters (QoS, lag-profiles, and so on) can be modified by the operator.

PXC sub-ports or LAGs can be removed from the FPE object only if they are not associated with an application.

2.7 LAG

Based on the IEEE 802.1ax standard (formerly 802.3ad), Link Aggregation Groups (LAGs) can be configured to increase the bandwidth available between two network devices, depending on the number of links installed. LAG also provides redundancy in the event that one or more links participating in the LAG fail. All physical links in a given LAG links combine to form one logical interface.

Packet sequencing must be maintained for any given session. The hashing algorithm deployed by the Nokia routers is based on the type of traffic transported to ensure that all traffic in a flow remains in sequence while providing effective load sharing across the links in the LAG.

LAGs must be statically configured or formed dynamically with Link Aggregation Control Protocol (LACP). The optional marker protocol described in IEEE 802.1ax is not implemented. LAGs can be configured on network and access ports.

The LAG load sharing is executed in hardware, which provides line rate forwarding for all port types.

The LAG implementation supports LAG with all member ports of the same speed and LAG with mixed port-speed members (see the sections that follow for details).

2.7.1 LACP

Under normal operation, all non-failing links in a given LAG will become active and traffic is load balanced across all active links. In some circumstances, however, this is not desirable. Instead, it is desired that only some of the links are active (for example, all links on the same IOM) and the other links be kept in stand-by condition.

LACP enhancements allow active LAG-member selection based on particular constraints. The mechanism is based on the IEEE 802.1ax standard so interoperability is ensured.

To use LACP on a given LAG, operator must enable LACP on the LAG including, if desired, selecting non-default LACP mode: active/passive and configuring administrative key to be used (**configure lag lacp**). In addition, an operator can configure a desired LACP transmit interval (**configure lag lacp-xmit-interval**).

When LACP is enabled, an operator can see LACP changes through traps and log messages logged against the LAG. See the TIMETRA-LAG-MIB.mib for more details.

2.7.1.1 LACP Multiplexing

The router supports two modes of multiplexing RX/TX control for LACP: coupled and independent.

In coupled mode (default), both RX and TX are enabled or disabled at the same time whenever a port is added or removed from a LAG group.

In independent mode, RX is first enabled when a link state is UP. LACP sends an indication to the far-end that it is ready to receive traffic. Upon the reception of this indication, the far-end system can enable TX. Therefore, in independent RX/TX control, LACP adds a link into a LAG only when it detects that the other end is ready to receive traffic. This minimizes traffic loss that might occur in coupled mode if a port is added into a LAG before notifying the far-end system or before the far-end system is ready to receive traffic. Similarly, on link removals from LAG, LACP turns off the distributing and collecting bit and informs the far-end about the state change. This allows the far-end side to stop sending traffic as soon as possible.

Independent control provides for lossless operation for unicast traffic in most scenarios when adding new members to a LAG or when removing members from a LAG. It also reduces loss for multicast and broadcast traffic.

Note that independent and coupled mode are interoperable (connected systems can have either mode set).

2.7.1.2 LACP Tunneling

LACP tunneling is supported on Epipe and VPLS services. In a VPLS, the Layer 2 control frames are sent out of all the SAPs configured in the VPLS. This feature should only be used when a VPLS emulates an end-to-end Epipe service (an Epipe configured using a three-point VPLS, with one access SAP and two access-uplink SAP/SDPs for redundant connectivity). The use of LACP tunneling is not recommended if the VPLS is used for multipoint connectivity. When a Layer 2 control frame is forwarded out of a dot1q SAP or a QinQ SAP, the SAP tags of the egress SAP are added to the packet.

The following SAPs can be configured for tunneling the untagged LACP frames (the corresponding protocol tunneling needs to be enabled on the port).

- If the port encapsulation is null, a null SAP can be configured on a port to tunnel these packets.
- If the port encapsulation is dot1q, either a dot1q explicit null SAP (for example, 1/1/10:0) or a dot1q default SAP (for example, 1/1/11:*) can be used to tunnel these packets.

- If the port encapsulation is QinQ, a 0.* SAP (for example, 1/1/10:0.*) can be used to tunnel these packets.

LAG port states may be impacted if LACP frames are lost due to incorrect prioritization and congestion in the network carrying the tunnel.

2.7.2 Active-Standby LAG Operation

An active-standby LAG provides redundancy by logically dividing LAG into subgroups. The LAG is divided into subgroups by either assigning each LAG's ports to an explicit subgroup (1 by default), or by automatically grouping all LAG's ports residing on the same line card into a unique sub-group (auto-iom) or by automatically grouping all LAG's ports residing on the same MDA into a unique sub-group (auto-mds). When a LAG is divided into sub-groups, only a single sub-group is elected as active. Which sub-group is selected depends on selection criterion chosen.

The active-standby decision for LAG member links is a local decision driven by preconfigured selection-criteria. When LACP is configured, this decision was communicated to remote system using LACP signaling.


To allow non-LACP operation, an operator must disable LACP on a given LAG and select transmitter-driven standby signaling (configure lag standby-signaling power-off). As a consequence, the transmit laser will be switched off for all LAG members in standby mode. On switch over (active-links failed) the laser will be switched on all standby LAG members so they can become active.

When the power-off is selected as the standby-signaling, the selection-criteria **best-port** can be used.

It is not possible to have an active LACP in power-off mode before the correct selection criteria is selected.


[Figure 37](#) shows how LAG in Active/Standby mode can be deployed towards a DSLAM access using sub-groups with auto-iom sub-group selection. LAG links are divided into two sub-groups (one per line card).

Figure 37 Active-Standby LAG Operation without Deployment Examples




In case of a link failure, as shown in [Figure 38](#) and [Figure 39](#), the switch over behavior ensures that all LAG-members connected to the same IOM as failing link become standby and LAG-members connected to other IOM become active. This way, QoS enforcement constraints are respected, while the maximum of available links is utilized.

Figure 38 LAG on Access Interconnection



OSSG095

Figure 39 LAG on Access Failure Switchover



OSSG096

2.7.3 LAG on Access QoS Consideration

The following section describes various QoS related features applicable to LAG on access.

2.7.3.1 Adapt QoS Modes

Link Aggregation is supported on the access side with access or hybrid ports. Similarly to LAG on the network side, LAG on access aggregates Ethernet ports into all active or active-standby LAG. The difference with LAG on networks lies in how the QoS or H-QoS is handled. Based on hashing configured, a given SAP's traffic can be sprayed on egress over multiple LAG ports or can always use a single port of a LAG. There are three user-selectable modes that allow operator to best adapt QoS configured to a LAG the SAPs are using:

1. adapt-qos distributed (default)

In a distributed mode the SLA is divided among all line cards proportionally to the number of ports that exist on that line card for a given LAG. For example a 100 Mb/s PIR with 2 LAG links on IOM A and 3 LAG links on IOM B would result in IOM A getting 40 Mb/s PIR and IOM B getting 60 Mb/s PIR. Because of this distribution, SLA can be enforced. The disadvantage is that a single flow is limited to IOM's share of the SLA. This mode of operation may also result in underrun due to a "hash error" (traffic not sprayed equally over each link). This mode is best suited for services that spray traffic over all links of a LAG.

2. adapt-qos link

In a link mode the SLA is given to each and every port of a LAG. With the example above, each port would get 100 Mb/s PIR. The advantage of this method is that a single flow can now achieve the full SLA. The disadvantage is that the overall SLA can be exceeded, if the flows span multiple ports. This mode is best suited for services that are guaranteed to hash to a single egress port.

3. adapt-qos port-fair

Port-fair distributes the SLA across multiple line cards relative to the number of active LAG ports per card (in a similar way to distribute mode) with all LAG QoS objects parented to scheduler instances at the physical port level (in a similar way to link mode). This provides a fair distribution of bandwidth between cards and ports whilst ensuring that the port bandwidth is not exceeded. Optimal LAG utilization relies on an even hash spraying of traffic to maximize the use of the schedulers' and ports' bandwidth. With the example above, enabling port-fair would result in all five ports getting 20 Mb/s.

When port-fair mode is enabled, per-Vport hashing is automatically disabled for subscriber traffic such that traffic sent to the Vport no longer uses the Vport as part of the hashing algorithm. Any QoS object for subscribers, and any QoS object for SAPs with explicitly configured hashing to a single egress LAG port, will be given the full bandwidth configured for each object (in a similar way to link mode). A Vport used together with an egress port scheduler is supported with a LAG in port-fair mode, whereas it is not supported with a distribute mode LAG.

4. adapt-qos distributed include-egr-hash-cfg

This mode can be considered a mix of link and distributed mode. The mode uses the configured hashing for LAG/SAP/service to choose either link or distributed adapt-qos modes. The mode allows:

- SLA enforcement for SAPs that through configuration are guaranteed to hash to a single egress link using full QoS per port (as per link mode)
- SLA enforcement for SAPs that hash to all LAG links proportional distribution of QoS SLA amongst the line cards (as per distributed mode)
- SLA enforcement for multi service sites (MSS) that contain any SAPs regardless of their hash configuration using proportional distribution of QoS SLA amongst the line cards (as per distributed mode)

The following caveats apply to adapt-qos distributed include-egr-hash-cfg:

- LAG mode must be access or hybrid.
- The operator cannot change from **adapt-qos distribute include-egr-hash-cfg** to **adapt-qos distribute** when link-map-profiles or per-link-hash is configured.
- The operator cannot change from **adapt-qos link** to **adapt-qos distribute include-egr-hash-cfg** on a LAG with any configuration.

[Table 35](#) shows examples of rate/BW distributions based on the **adapt-qos** mode used.

Table 35 Adapt QoS Bandwidth/Rate Distribution

	distribute	link	port-fair	distribute include-egr-hash-cfg
SAP Queues	% # local links ¹	100% rate	100% rate (SAP hash to one link) or %# all links ² (SAP hash to all links)	100% rate (SAP hash to one link) or % # local linksa (SAP hash to all links)

Table 35 **Adapt QoS Bandwidth/Rate Distribution (Continued)**

	distribute	link	port-fair	distribute include-egr-hash-cfg
SAP Scheduler	% # local linksa	100% bandwidth	100% rate (SAP hash to one link) or %# all linksb (SAP hash to all links)	100% bandwidth (SAP hash to a one link) or % # local linksa (SAP hash to all links)
SAP MSS Scheduler	% # local linksa	100% bandwidth	% # local linksa	% # local linksa

Notes:

1. * % # local links = $X * (\text{number of local LAG members on a given line card} / \text{total number of LAG members})$
2. %# all links = $X * (\text{link speed}) / (\text{total LAG speed})$

2.7.3.2 Per-fp-ing-queuing

Per-fp-ing-queuing optimization for LAG ports provides the ability to reduce the number of hardware queues assigned on each LAG SAP on ingress when the flag at LAG level is set for per-fp-ing-queuing.

When the feature is enabled in the **config>lag>access** context, the queue allocation for SAPs on a LAG will be optimized and only one queuing set per ingress forwarding path (FP) is allocated instead of one per port.

The following rules will apply for configuring the per-fp-ing-queuing at LAG level:

- To enable per-fp-ing-queuing, the LAG must be in access mode
- The LAG mode cannot be set to network mode when the feature is enabled
- Per-fp-ing-queuing can only be set if no port members exists in the LAG
- Per-fp-ing-queuing cannot be set if LAG's port-type is hsmdda

2.7.3.3 Per-fp-egr-queuing

Per-fp-egr-queuing optimization for LAG ports provides the ability to reduce the number of egress resources consumed by each SAP on a LAG, and by any encap groups that exist on those SAPs.

When the feature is enabled in the **config>lag>access** context, the queue and virtual scheduler allocation will be optimized. Only one queuing set and one H-QoS virtual scheduler tree per SAP/encap group will be allocated per egress forwarding path (FP) instead of one set per each port of the LAG. In case of a link failure/recovery, egress traffic uses failover queues while the queues are moved over to a newly active link.

Per-fp-egr-queuing can be enabled on existing LAG with services as long as the following conditions are met.

- The LAG's mode must be **access** or **hybrid**.
- The LAG's port-type must be **standard**.
- The LAG must have either **per-link-hash** enabled or all SAPs on the LAG must use **per-service-hashing** only and be of a type: VPLS SAP, i-VPLS SAP, or e-Pipe VLL or PBB SAP.

To disable per-fp-egr-queuing, all ports must first be removed from a given LAG.

2.7.3.4 Per-fp-sap-instance

Per-fp-sap-instance optimization for LAG ports provides the ability to reduce the number of SAP instance resources consumed by each SAP on a lag.

When the feature is enabled, in the **config>lag>access** context, a single SAP instance is allocated on ingress and on egress per each forwarding path instead of one per port. Thanks to an optimized resource allocation, the SAP scale on a line card will increase, if a LAG has more than one port on that line card. Because SAP instances are only allocated per forwarding path complex, hardware reprogramming must take place when as result of LAG links going down or up, a SAP is moved from one LAG port on a given line card to another port on a given line card within the same forwarding complex. This results in an increased data outage when compared to per-fp-sap-instance feature being disabled. During the reprogramming, failover queues are used when SAP queues are reprogrammed to a new port. Any traffic using failover queues will not be accounted for in SAPs statistics and will be processed at best-effort priority.

The following rules apply when configuring a per-fp-sap-instance on a given LAG:

- Per-fp-sap-ing-queuing and per-fp-sap-egr-queuing must be enabled.
- The functionality can be enabled/disabled on LAG with no member ports only. Services can be configured.

Other caveats:

- SAP instance optimization applies to LAG-level. Whether a LAG is sub-divided into sub-groups or not, the resources are allocated per forwarding path for all complexes LAG's links are configured on (i.e. irrespective of whether a given sub-group a SAP is configured on uses that complex or not).
- Egress statistics continue to be returned per port when SAP instance optimization is enabled. If a LAG links are on a single forwarding complex, all ports but one will have no change in statistics for the last interval – unless a SAP moved between ports during the interval.
- Rollback that changes per-fp-sap-instance configuration is service impacting.

2.7.4 LAG and ECMP Hashing

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and/or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a given service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred to as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a given flow at the same time, only one method will be selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

2.7.4.1 Per Flow Hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any given flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not desired, other hashing methods outlined in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and/or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following outlines default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
 - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
 - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data will be hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing will only be performed twice to get the LAG port.
 - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.

- The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.
- MPLS packet hashing at an LSR is based on the whole label stack, along with the incoming port and system IP address. Note that the EXP/TTL information in each label is not included in the hash algorithm. This method is referred to as *Label-Only Hash* option and is enabled by default, or can be re-instated in CLI by entering the `lbl-only` option. A few options to further hash on the headers in the payload of the MPLS packet are also provided. For more details, see [Changing Default Per Flow Hashing Inputs](#).
- VLL traffic from a service access point is not sprayed on a per-packet basis, but as for VPLS flooded traffic, the service ID selects one of the ECMP/LAG paths. The exception to this is when shared-queuing is configured on an e-pipe SAP, i-pipe SAP, or f-pipe SAP, or when H-POL is configured on an e-pipe SAP. In those cases, traffic spraying is the same as for VPLS known unicast traffic. Packets of the above VLL services received on a spoke-SDP are sprayed the same as for VPLS known unicast traffic.
- Note that a-pipe and c-pipe VLL packets are always sprayed based on the service-id in both directions.
- Multicast IP traffic is hashed based on an internal multicast ID, which is unique for every record similar to VPLS multicast traffic with IGMP snooping enabled.

If the ECMP index results in the selection of a LAG as the next hop, then the hash result is hashed again and the result of the second hash is input to the modulo like operation to determine the LAG port selection.

When the ECMP set includes an IP interface configured on a spoke-SDP (IES/VP RN spoke interface), or a Routed VPLS spoke-sdp interface, the unicast IP packets—which will be sprayed over this interface—will not be further sprayed over multiple RSVP LSPs/LDP FEC (part of the same SDP), or GRE SDP ECMP paths. In this case, a single RSVP LSP, LDP FEC next-hop or GRE SDP ECMP path will be selected based on a modulo operation of the service ID. In case the ECMP path selected is a LAG, the second round of the hash, hashes traffic based on the system, port or interface load-balancing settings.

In addition to the above outlined per-flow hashing inputs, the system supports multiple options to modify default hash inputs.

2.7.4.1.1 Changing Default Per Flow Hashing Inputs

For some traffic patterns or specific deployments, per-flow hashing is desired but the hashing result using default hash inputs as outlined above may not produce the desired distribution. To alleviate this issue, the system allows operators to modify the load balancing algorithm at the system, interface or port level.

LSR Hashing

By default, the LSR hash routine operates on the label stack only. However, there is also the ability to hash on the IP header if a packet is IP. An LSR will consider a packet to be IP if the first nibble following the bottom of the label stack is either 4 (IPv4) or 6 (IPv6). This allows the user to include an IP header in the hashing routine at an LSR for the purpose of spraying labeled IP packets over multiple equal cost paths in ECMP in an LSP and/or over multiple links of a LAG group in all types of LSPs.

The user enables the LSR hashing on label stack and/or IP header by entering the following system-wide command: **config>system>load-balancing>lsr-load-balancing [lbl-only | lbl-ip | ip-only]**

By default, the LSR falls back to the hashing on label stack only. This option is referred to as lbl-only and the user can revert to this behavior by entering one of the two commands:

config>system>load-balancing>lsr-load-balancing lbl-only

config>system>load-balancing>no lsr-load-balancing

The user can also selectively enable or disable the inclusion of label stack and IP header in the LSR hash routine for packets received on a specific network interface by entering the following command:

config>router>if>load-balancing>lsr-load-balancing [lbl-only | lbl-ip | ip-only | eth-encap-ip | lbl-ip-l4-teid]

This provides some control to the user such that this feature is disabled if labeled packets received on a specific interface include non IP packets that can be confused by the hash routine for IP packets. These could be VLL and VPLS packets without a PW control word.

When the user performs the **no** form of this command on an interface, the interface inherits the system level configuration.

LSR Default Hash Routine—Label-Only Hash Option

The following is the behavior of ECMP and LAG hashing at an LSR in the existing implementation. These are performed in two rounds.

First the ECMP hash. It consists of an initial hash based on the source port/system IP address. Each label in the stack is then hashed separately with the result of the previous hash, up to a maximum of 16 labels. The net result will be used to select which LSP next-hop to send the packet to using a modulo operation of the net result with the number of next-hops.

This same net result will feed to a second round of hashing if there is LAG on the egress port where the selected LSP has its NHLFE programmed.

LSR Label-IP Hash Option Enabled

In the first hash round for ECMP, the algorithm will parse down the label stack and once it hits the bottom it checks the next nibble. If the nibble value is 4 then it will assume it is an IPv4 packet. If the nibble value is 6 then it will assume it is an IPv6 packet. In both cases, the result of the label hash is fed into another hash along with source and destination address fields in the IP packet header. Otherwise, it will just use the label stack hash already calculated for the ECMP path selection.

If there are more than five labels in the stack, then the algorithm will also use the result of the label hash for the ECMP path selection.

The second round of hashing for LAG re-uses the net result of the first round of hashing. This means IPv6 packets will continue to be hashed on label stack only.

LSR IP-Only Hash Option Enabled

This option behaves like the label-IP hash option except that when the algorithm reached the bottom of the label stack in the ECMP round and finds an IP packet, it throws the outcome of the label hash and only uses the source and destination address fields in the IP packet's header.

LSR Ethernet Encapsulated IP Hash only Option Enabled

This option behaves like LSR IP only hash except for how the IP SA/DA information is found. The following conditions are verified to find IP SA/DA for hash.

- Label stack must not exceed 3 labels deep

- After the bottom of the stack is reached, the hash algorithm verifies that what follows is Ethernet II untagged frame (by looking at the value of ethertype at the expected packet location whether it contains Ethernet encapsulated IPv4 (0x0800) or IPv6 (0x86DD) value.

When the ethertype verification passes, the first nibble of the expected IP packet location is then verified to be 4 (IPv4) or 6 (IPv6).

LSR Hashing of MPLS-over-GRE Encapsulated Packet

When the router removes the GRE encapsulation, pops one or more labels and then swaps a label, it acts as an LSR. The LSR hashing for packets of a MPLS-over-GRE SDP or tunnel follows a different procedure which is enabled automatically and overrides the LSR hashing option enabled on the incoming network IP interface (**lsr-load-balancing {lbl-only | lbl-ip | ip-only | eth-encap-ip | lbl-ip-l4-teid}**).

On a packet by packet basis, the new hash routine will parse through the label stack and:

- a. hash on the SA/DA fields and the Layer 4 SRC/DST Port fields of the inner IPv4/IPv6 header.
- b. If the GRE header and label stack sizes are such that the Layer4 SRC/DST Port fields are not read, it will hash on the SA/DA fields of the inner IPv4/IPv6 header.
- c. If the GRE header and label stack sizes are such that the SA/DA fields of the inner IPv4/IPv6 header are not read, it will hash on the SA/DA fields of the outer IPv4 header.

LSR Hashing when an Entropy Label is Present in the Packet's Label Stack

The LSR hashing procedures are modified as follows:

- If the **lbl-only** hashing option is enabled, or if one of the other LSR hashing options are enabled but an IPv4 or IPv6 header is not detected below the bottom of the label stack, the LSR hashes on the Entropy Label (EL) only.
- If the **lbl-ip** option is enabled, the LSR hashes on the EL and the IP headers.
- If the **ip-only** or **eth-encap-ip** is enabled, the LSR hashes on the IP headers only.

Layer 4 Load Balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be subdivided into multiple finer-granularity flows if the ports used between a given SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP Load Balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths will have a lower chance of always using the same path to a given destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic will use the same nexthop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XP/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

TEID Hash for GTP-Encapsulated Traffic

This options enables TEID hashing on Layer 3 interfaces. The hash algorithm identifies GTP-C or GTP-U by looking at the UDP destination port (2123 or 2152) of an IP packet to be hashed. If the value of the port matches, the packet is assumed to be GTP-U/C. For GTPv1 packets TEID value from the expected header location is then included in hash. For GTPv2 packets the TEID flag value in the expected header is additionally checked to verify whether TEID is present. If TEID is present, it is included in hash algorithm inputs. TEID is used in addition to GTP tunnel IP hash inputs: SA/DA and SPort/DPort (if Layer 4 load balancing is enabled). If a non-GTP packet is received on the GTP UDP ports above, the packets will be hashed as GTP.

Source-Only/Destination-Only Hash Inputs

This option allows an operator to only include source parameters or only include destination parameters in the hash for inputs that have source/destination context (such as IP address and Layer 4 port). Parameters that do not have source/destination context (such as TEID or System IP for example) are also included in hash as per applicable hash configuration. The functionality allows, among others, to ensure that both upstream and downstream traffic hash to the same ECMP path/LAG port on system egress when traffic is sent to a hair-pinned appliance (by configuring source-only hash for incoming traffic on upstream interfaces and destination-only hash for incoming traffic on downstream interfaces).

Enhanced Multicast Load Balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and/or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the [LSR Hashing](#) section.



Note: Enhanced multicast load balancing is not supported with Layer 2 and ESM services. It is supported on all platforms except for the 7750 SR-c4 and SR-c12 and the 7450 ESS in standard mode.

Security Parameter Index (SPI) Load Balancing

IPSec tunneled traffic transported over LAG typically falls back to IP header hashing only. For example, in LTE deployments, TEID hashing cannot be performed because of encryption, and the system performs IP-only tunnel-level hashing. Because each SPI in the IPSec header identifies a unique SA, and thus flow, these flows can be hashed individually without impacting packet ordering. In this way, SPI load balancing provides a mechanism to improve the hashing performance of IPSec encrypted traffic.

The system allows enabling SPI hashing per Layer 3 interface (this is the incoming interface for hash on system egress)/Layer 2 VPLS service. When enabled, an SPI value from ESP/AH header is used in addition to any other IP hash input based on per-flow hash configuration: source/destination IPv6 addresses, Layer 4 source/dest ports in case NAT traversal is required (Layer 4 load-balancing is enabled). If the ESP/AH header is not present in a packet received on a given interface, the SPI will not be part of the hash inputs, and the packet is hashed as per other hashing configurations. SPI hashing is not used for fragmented traffic to ensure first and subsequent fragments use the same hash inputs.

SPI hashing is supported for IPv4 and IPv6 tunnel unicast traffic and for multicast traffic (mc-enh-load-balancing must be enabled) on all platforms and requires Layer 3 interfaces or VPLS service interfaces with SPI hashing enabled to reside on IOM3-XP or newer line-cards.

2.7.4.2 Per Link Hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a given SAP or network interface uses a single LAG port on egress. Because all traffic for a given SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a given LAG have statistically similar BW requirements (since per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

2.7.4.2.1 Weighted per-link-hash

Weighted per-link-hash allows higher control in distribution of SAPs/interfaces/subscribers across LAG links when significant differences in SAPs/interfaces/subscribers bandwidth requirements could lead to an unbalanced distribution bandwidth utilization over LAG egress. The feature allows operators to configure for each SAPs/interfaces/subscribers on a LAG one of three unique classes and a weight value to be used to when hashing this service/subscriber across the LAG links. SAPs/interfaces/subscribers are hashed to LAG links, such that within each class the total weight of all SAPs/interfaces/subscribers on each LAG link is as close as possible to each other.

Multiple classes allow grouping of SAPs/interfaces/subscribers by similar bandwidth class/type. For example a class can represent: voice – negligible bandwidth, Broadband – 10 to 100 Mb/s, Extreme Broadband – 300 Mb/s and above types of service. If a class and weight are not specified for a given service or subscriber, values of 1 and 1 are used respectively.

The following algorithm hashes SAPs, interfaces, and subscribers to LAG egress links:

- TPSDA subscribers are hashed to a LAG link when subscribers are active, MSE SAPs/interfaces are hashed to a LAG link when configured
- For a new SAP/interface/subscriber to be hashed to an egress LAG link:
 - Select active link with the smallest current weight for the SAP/network/subscriber class
- On a LAG link failure:
 - Only SAPs/interfaces/subscribers on a failed link are rehashed over the remaining active links
 - Processing order: Per class from lowest numerical, within each class per weight from highest numerical value
- LAG link recovery/new link added to a LAG:
 - auto-rebalance disabled: Existing SAPs/interfaces/subscribers remain on the currently active links, new SAPs/interfaces/subscribers naturally prefer the new link until balance reached.
 - auto-rebalance is enabled: When a new port is added to a LAG a non-configurable 5 second rebalance timer is started. Upon timer expiry, all existing SAPs/interfaces/subscribers are rebalanced across all active LAG links minimizing the number of SAPs/interfaces/subscribers moved to achieve rebalance. The rebalance timer is restarted if a new link is added while the timer is running. If a port bounces 5 times within a 5 second interval, the port is quarantined for 10 seconds. This behavior is not configurable.

- On a LAG start up, the rebalance timer is always started irrespective of auto-rebalance configuration to avoid hashing SAPs/interfaces/subscribers to a LAG before ports have a chance to come UP.
- Weights for network interfaces are separated from weights for access SAPs/interfaces/subscribers.
- On a mixed-speed LAG, link selection is made with link speeds factoring into the overall weight for the same class of traffic. This means that higher-speed links will be preferred over lower-speed links.

Optionally an operator can use a **tools perform lag load-balance** command to manually re-balance all weighted per-link-hashed SAPs/interfaces/subscribers on a LAG. The rebalance follows the algorithm as used on a link failure moving SAPs/interfaces/subscribers to different LAG links to minimize SAPs/interfaces/subscribers impacted.

Along with the caveats for standard per-link hashing, the following caveats exist:

- When weighted per-link-hash is deployed on a given LAG, no other methods of hash for subscribers/SAPs/interfaces on that LAG (like service hash or LAG link map profile) should be deployed, since the weighted hash is not able to account for loads placed on LAG links by subscriber/SAPs/interfaces using the other hash methods.
- For the TPSDA model only the 1:1 (subscriber to SAP) model is supported.

This feature will not operate properly if the above conditions are not met.

2.7.4.3 Explicit Per Link Hash Using LAG Link Mapping Profiles

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. LAG link mapping profile feature gives operators full control of which links SAPs/network interface use on a LAG egress and how the traffic is rehashed on a LAG link failure. Some benefits that such functionality provides include:

- Ability to perform management level admission control onto LAG ports thus increasing overall LAG BW utilization and controlling LAG behavior on a port failure.
- Ability to strictly enforce QoS contract on egress for a SAP/network interface or a group of SAPs/network interfaces by forcing it/them to egress over a single port and using **access adapt-qos** link or port-fair mode.

To enable LAG Link Mapping Profile Feature on a given LAG, operators configure one or more of the available LAG link mapping profiles on the LAG and then assign that profile(s) to all or a subset of SAPs and network interfaces as needed. Enabling per LAG link Mapping Profile is allowed on a LAG with services configured, a small outage may take place as result of re-hashing SAP/network interface when a lag profile is assigned to it.

Each LAG link mapping profile allows operators to configure:

- Primary link—defines a port of the LAG to be used by a SAP/network interface when the port is UP. Note that a port cannot be removed from a LAG if it is part of any LAG link profile.
- Secondary link—defines a port of the LAG to be used by a SAP/network interface as a backup when the primary link is not available (not configured or down) and the secondary link is UP.
- Mode of operation when neither primary, nor secondary links are available (not configured or down):
 - **discard** – traffic for a given SAP/network interface will be dropped to protect other SAPs/network interfaces from being impacted by re-hashing these SAPs/network interfaces over remaining active LAG ports.



Note: SAP/network interface status will not be affected when primary and secondary links are unavailable, unless an OAM mechanism that follows the data path hashing on egress is used and will cause a SAP/network interface to go down.

- **per-link-hash** – traffic for a given SAP/network interface will be re-hashed over remaining active ports of a LAG links using per-link-hashing algorithm. This behavior ensures SAP/network interfaces using this profile will be given available resources of other active LAG ports even if that means impacting other SAP/network interfaces on the LAG. The system will use the QoS configuration to provide fairness and priority if congestion is caused by the default-hash recovery.

LAG link mapping profiles, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid)
- All ports of a LAG on a given router must belong to a single sub-group.
- Access adapt-qos mode is distribute include-egr-hash-cfg.

LAG link mapping profile can co-exist with any-other hashing used over a given LAG (for example, per flow hashing or per-link-hashing). SAPs/network interfaces that have no link mapping profile configured will be subject to LAG hashing, while SAPs/network interfaces that have configured LAG profile assigned will be subject to LAG link mapping behavior, which is described above.

2.7.4.4 Consistent Per Service Hashing

The hashing feature described in this section applies to traffic going over LAG, Ethernet tunnels (eth-tunnel) in loadsharing mode, or CCAG load balancing for VSM redundancy. The feature does not apply to ECMP.

Per-service-hashing was introduced to ensure consistent forwarding of packets belonging to one service. The feature can be enabled using the **[no] per-service-hashing** configuration option under **config>service>epipe>load-balancing** and **config>service>vpls>load-balancing**, valid for Epipe, VPLS, PBB Epipe, IVPLS and BVPLS.

The following behavior applies to the usage of the [no] per-service-hashing option.

- The setting of the PBB Epipe/I-VPLS children dictates the hashing behavior of the traffic destined to or sourced from an Epipe/I-VPLS endpoint (PW/SAP).
- The setting of the B-VPLS parent dictates the hashing behavior only for transit traffic through the B-VPLS instance (not destined to or sourced from a local I-VPLS/Epipe children).

The following algorithm describes the hash-key used for hashing when the new option is enabled:

- If the packet is PBB encapsulated (contains an I-TAG ethertype) at the ingress side and enters a B-VPLS service, use the ISID value from the I-TAG. For PBB encapsulated traffic entering other service types, use the related service ID.
- If the packet is not PBB encapsulated at the ingress side
 - For regular (non-PBB) VPLS and EPIPE services, use the related service ID
 - If the packet is originated from an ingress IVPLS or PBB Epipe SAP
 - If there is an ISID configured use the related ISID value
 - If there is no ISID configured use the related service ID
 - For BVPLS transit traffic use the related flood list id
 - Transit traffic is the traffic going between BVPLS endpoints
 - An example of non-PBB transit traffic in BVPLS is the OAM traffic
- The above rules apply regardless of traffic type

- Unicast, BUM flooded without MMRP or with MMRP, IGMP snooped

Operators may sometimes require the capability to query the system for the link in a LAG or Ethernet tunnel that is currently assigned to a given service-id or ISID. This capability is provided using the **tools>dump>map-to-phy-port {ccag ccag-id | lag lag-id | eth-tunnel tunnel-index} {isid isid [end-isid isid] | service servid-id | svc-name [end-service service-id | svc-name]} [summary]** command.

A sample usage is as follows:

```
A:Dut-B# tools dump map-to-phy-port lag 11 service 1
```

ServiceId	ServiceName	ServiceType	Hashing	Physical Link
1		i-vpls	per-service(if enabled)	3/2/8

```
A:Dut-B# tools dump map-to-phy-port lag 11 isid 1
```

ISID	Hashing	Physical Link
1	per-service(if enabled)	3/2/8

```
A:Dut-B# tools dump map-to-phy-port lag 11 isid 1 end-isid 4
```

ISID	Hashing	Physical Link
1	per-service(if enabled)	3/2/8
2	per-service(if enabled)	3/2/7
3	per-service(if enabled)	1/2/2
4	per-service(if enabled)	1/2/3

2.7.4.5 ESM – LAG Hashing per Vport

2.7.4.5.1 Background

Vport is a router BNG representation of a remote traffic aggregation point in the access network. It is a level in the hierarchical QoS model implemented within the BNG that requires QoS treatment.

When the BNG is connected to access network via LAG, a VPort construct within the BNG is instantiated per member link on that LAG. Each instance of the Vport in such a configuration receives the entire amount of configured bandwidth. When traffic is sprayed in a per-subscriber fashion over member links in an LAG without awareness of the Vport, it can lead to packet drops on one member link irrespective of the relative traffic priority on another LAG member link in the same Vport. The reason is that multiple Vport instances of the same Vport on different LAG member links are not aware of each other.

With a small number of subscribers per Vport and a great variation in bandwidth service offering per subscriber (from Mb/s to Gb/s), there is a great chance that the load distribution between the member links will be heavily unbalanced. For example, if the lag consists of two member links on the same IOM, three 1Gb/s high priority subscribers can saturate the 2 Gb/s Vport bandwidth on one member link of the LAG. And all the while, twenty low priority 10 Mb/s subscribers that are using the other link are significantly under-utilizing available bandwidth on the corresponding Vport.

To remedy this situation, all traffic flowing through the same Vport must be hashed to a single LAG member link. This way, the traffic treatment will be controlled by a single Vport instance, and achieve a desired behavior where low priority 10 Mb/s subscribers traffic will be affected before any traffic from the high priority subscribers.

2.7.4.5.2 Hashing per Vport

Hashing traffic per Vport ensures that the traffic on the same PON (or DSLAM) traverse the same Vport, and therefore, it is the same member link that this Vport is associated with. The Vport instances of the same Vport on another member links are irrelevant for QoS treatment.

The Vport for Nokia routers is referenced via inter-dest-string, which can be returned via RADIUS. For this reason, the terms hashing per inter-dest-string or hashing per Vport can be interchangeably used.

If the subscriber is associated with a Vport, hashing will be automatically performed per inter-dest-string. In case that no such association exists, hashing will default to per-subscriber hashing.

In certain cases, S-vlan tag can represent Vport. In such a case, per S-vlan hashing is desired. This can be implicitly achieved by the following configuration:

```
configure
  subscr-mgmt
    msap-policy <name>
    sub-sla-mgmt
    def-inter-dest-id use-top-queue

configure
  port <port-id>
    ethernet
      access
      egress
    vport <name>
      host-match dest <s-tag>
```

Through this CLI hierarchy, S-tag is implicitly associated with the inter-dest-string and consequently with the Vport.



Note: LAG hashing parameters that are configured under **config>lag**, for example, **per-link-hash**, take precedence and are incompatible with the **vport-hashing** command.

2.7.4.5.3 Vport Hashing over Different Forwarding Complexes

This feature enables the use of vPort ID as the hashing key, enabling an active-active LAG configuration to span more than one forwarding complex. The feature does not interoperate with AA capabilities tied to a specific subscriber.

When using Vport hashing, **adapt-qos link** mode is recommended on the access interface.

This feature can be enabled using the CLI command **config>subscr-mgmt>sub-profile>vport-hashing**. On the HSM DA, the secondary shaper supports Vport hashing through the **config>subscr-mgmt>sub-profile>secondary-shaper-hashing** command.

2.7.4.5.4 Multicast Consideration

Multicast traffic that is directly replicated per subscriber follows the same hashing algorithm as the rest of the subscribers (per inter-dest-string hashing).

Multicast traffic that is redirected to a regular Layer 3 interface outside of the ESM will be hashed per destination group (or IP address).

2.7.4.5.5 VPLS and Capture SAP Considerations

VPLS environment in conjunction with ESM allows hashing based on destination mac address. This is achieved through the following CLI hierarchy:

```
configure
  service vpls <vpls-id>
    sap lag-<id>
  sub-sla-mgmt
    mac-da-hashing
```

Note that this is only applicable to Layer 2 ESM. In the case where this is configured and Vport hashing is required, the following order of evaluation must be executed:

1. Hashing based on subscriber-id or inter-dest-string
2. If configured, mac-da-hashing

Hashing per inter-dest-string will win if a <Vport, subscriber> association is available at the same time as the mac-da-hashing is configured.

The Mac-da-hashing mechanism cannot transition from a capture SAP to a derived MSAP.

2.7.5 LAG Hold Down Timers

Operators can configure multiple hold down timers that allow control how quickly LAG responds to operational port state changes. The following timers are supported:

1. Port-level hold-time up/down timer
This optional timer allows operator to control delay for adding/removing a port from LAG when the port comes UP/goes DOWN. Each LAG port runs the same value of the timer, configured on the primary LAG link. See the Port Link Dampening description in [Port Features](#) for more details on this timer.
2. Sub-group-level hold-time timer
This optional timer allows operator to control delay for a switch to a new candidate sub-group selected by LAG sub-group selection algorithm from the current, operationally UP sub-group. The timer can also be configured to never expire, which prevents a switch from operationally up sub-group to a new candidate sub-group (manual switchover is possible using tools perform force lag command). Note that, if the port link dampening is deployed, the port level timer must expire before the sub-group-selection takes place and this timer is started. Sub-group-level hold-down timer is supported with LAGs running LACP only.
3. LAG-level hold-time down timer
This optional timer allows operator to control delay for declaring a LAG operationally down when the available links fall below the required port/BW minimum. The timer is recommended for LAG connecting to MC-LAG systems. The timer prevents a LAG going down when MC-LAG switchover executes break-before-make switch. Note that, if the port link dampening is deployed, the port level timer must expire before the LAG operational status is processed and this timer is started.

2.7.6 BFD over LAG Links

The router supports the application of BFD to monitor individual LAG link members to speed up the detection of link failures. When BFD is associated with an Ethernet LAG, BFD sessions are setup over each link member, and are referred to as micro-BFD sessions. A link is not operational in the associated LAG until the associated micro-BFD session is fully established. In addition, the link member is removed from the operational state in the LAG if the BFD session fails.

When configuring the local and remote IP address for the BFD over LAG link sessions, the **local-ip** parameter should always match an IP address associated with the IP interface to which this LAG is bound. In addition, the **remote-ip** parameter should match an IP address on the remote system and should also be in the same subnet as the **local-ip** address. If the LAG bundle is re-associated with a different IP interface, the **local-ip** and **remote-ip** parameters should be modified to match the new IP subnet. The **local-ip** and **remote-ip** values do not have to match in the case of hybrid mode, q-tag or QinQ tagging.

2.7.7 Mixed Port-Speed LAG Support

Nokia routers support mixing different speed member ports in a single LAG. The LAG must be configured explicitly to allow mixed port-speed operation through the port-weight-speed command. The port-weight-speed defines both the lowest port speed for a member port in that LAG and the type of higher speed ports allowed to be mixed in the same LAG. For example, port-weight-speed 10 defines the minimum member port speed of 10GE and allows addition of any port that has a speed, which is a multiple of 10GE as long as the mix is supported by a given release, refer to specific Release Notes. Any LAG can be configured to support mixed port-speed operation.

For mixed port-speed LAGs:

- Both LACP and non-LACP configurations are supported. With LACP enabled, LACP is unaware of physical port differences.
- QoS is distributed proportionally to port-speed, unless explicitly configured not to do so (see **internal-scheduler-weight-mode**).
- User data traffic is hashed proportionally to port speed when any per-flow hash is deployed.
- CPM-originated OAM control traffic that requires per LAG hashing is hashed per physical port.

- Nokia recommends that operators use **weight-threshold** instead of **port-threshold** to control LAG operational status. For example, when 10GE and 100GE ports are mixed in a LAG, each 10GE port will have a weight of 1, while each 100GE port will have a weight of 10.

The **weight-threshold** command can also be used for LAGs with all ports of equal speed to allow a common operational model. For example, each port has a weight of 1 to mimic **port-threshold** and its related configuration.

- Nokia recommends that operators use weight-based thresholds for other system configurations that react to operational change of LAG member ports, like MCAC (see **use-lag-port-weight**) and VRRP (see **weight-down**).
- When sub-groups are used, the following behavior should be noted for selection criteria:
 - highest-count – continues to operate on physical link counts. Therefore, a sub-group with lower speed links will be selected even if its total bandwidth is lower. For example: a 4 * 10GE subgroup will be selected over a 100GE + 1 GE sub-group).
 - highest-weight – continues to operate on operator-configured priorities. Therefore, it is expected that configured weights take into account the proportional bandwidth difference between member ports to achieve the desired behavior. For example, to favor sub-groups with higher bandwidth capacity but lower link count in a 10GE/100GE LAG, 100GE ports need to have their priority set to a value that is at least 10 times that of the 10GE ports priority value.
 - best-port – continues to operate on operator-configured priorities. Therefore, it is expected that the configured weights will take into account proportional bandwidth difference between member ports to achieve the desired behavior.

Operators can add higher speed member ports to an existing LAG in service when all ports of the LAG have the speed as selected by port-weight-speed or when port-weight-speed is disabled (non-mixed port-speed operation). To do so, first port-based thresholds related to that LAG should be switched to weight-based thresholds, and then port-speed-weight should be set to the port speed of the existing member ports. After that, operators can add higher speed ports adjusting weight-based thresholds as required.

Similarly, operators can disable mixed port-speed operation in service if all ports have the same port speed and port-weight-speed equals member ports' speed. Weight-based thresholds may remain in use for the LAG.

Feature limitations:

- supported as of Release 14.0R1
- supported on network, access, and hybrid mode LAGs, including MC-LAG

- supported for standard-port LAGs and on 10GE WAN/100GE LAN port combinations
- PIM lag-usage-optimization is not supported and must not be configured
- LAG member links must have the default configuration for **config port ethernet egress-rate/ingress-rate**
- not supported on the 7450 ESS-6V
- not supported for ESM


2.7.7.1 LAG Upgrade

Migrating LAGs to higher speed links involves using mixed-speed LAGs initially, and later removing lower speed links. However, a consequence is that the lower speed links in the mixed-speed LAG set the member link limit. Even after all lower speed links are removed, the higher-speed links maintain a higher weight and this limits how many physical links that a mixed-port speed LAG can include.

LAG upgrade support allows migration from 1GE to 10GE to 40/100GE without removing all the ports from the LAG.

LAG upgrade support requires turning on mixed-speed LAG and adding higher speed links to an existing LAG. Once the lower speed links are removed, the **no port-weight-speed** command turns off mixed-speed LAG and to re-calibrate the number of logical links. [Figure 40](#) illustrates the steps in this scenario.

Figure 40 LAG Upgrade (Mixed Speed LAGs)



0949

If a 10GE or 100GE port is allocated as 10 links, it would be converted to one link per port if all the ports in the LAG are the same speed.

2.7.8 Multi-Chassis LAG

This section describes the Multi-Chassis LAG (MC-LAG) concept. MC-LAG is an extension of a LAG concept that provides node-level redundancy in addition to link-level redundancy provided by “regular LAG”.

Typically, MC-LAG is deployed in a network-wide scenario providing redundant connection between different end points. The whole scenario is then built by combination of different mechanisms (for example, MC-LAG and redundant pseudowire to provide e2e redundant p2p connection or dual homing of DSLAMs in Layer 2/3 TPSDA).

2.7.8.1 Overview

Multi-chassis LAG is a method of providing redundant Layer 2/3 access connectivity that extends beyond link level protection by allowing two systems to share a common LAG end point.

The multi-service access node (MSAN) node is connected with multiple links towards a redundant pair of Layer 2/3 aggregation nodes such that both link and node level redundancy, are provided. By using a multi-chassis LAG protocol, the paired Layer 2/3 aggregation nodes (referred to as redundant-pair) appears to be a single node utilizing LACP towards the access node. The multi-chassis LAG protocol between a redundant-pair ensures a synchronized forwarding plane to and from the access node and synchronizes the link state information between the redundant-pair nodes such that proper LACP messaging is provided to the access node from both redundant-pair nodes.

In order to ensure SLAs and deterministic forwarding characteristics between the access and the redundant-pair node, the multi-chassis LAG function provides an active/standby operation to and from the access node. LACP is used to manage the available LAG links into active and standby states such that only links from 1 aggregation node are active at a time to/from the access node.

Alternatively, when access nodes do not support LACP, the **power-off** option can be used to enforce the active/standby operation. In this case, the standby ports are **trx_disabled** (power off transmitter) to prevent usage of the LAG member by the access-node. Characteristics related to MC are:

- Selection of the common system ID, system-priority and administrative-key are used in LACP messages so partner systems consider all links as the part of the same LAG.
- Extension of selection algorithm in order to allow selection of active sub-group.

- The sub-group definition in LAG context is still local to the single box, meaning that even if sub-groups configured on two different systems have the same sub-group-id they are still considered as two separate subgroups within a specified LAG.
- Multiple sub-groups per PE in an MC-LAG is supported.
- In case there is a tie in the selection algorithm, for example, two sub-groups with identical aggregate weight (or number of active links) the group which is local to the system with lower system LACP priority and LAG system ID is taken.
- Providing inter-chassis communication channel allows inter-chassis communication to support LACP on both system. This communication channel enables the following:
 - Supports connections at the IP level which do not require a direct link between two nodes. The IP address configured at the neighbor system is one of the addresses of the system (interface or loop-back IP address).
 - The communication protocol provides heartbeat mechanism to enhance robustness of the MC-LAG operation and detecting node failures.
 - Support for operator actions on any node that force an operational change.
 - The LAG group-ids do not have to match between neighbor systems. At the same time, there can be multiple LAG groups between the same pair of neighbors.
 - Verification that the physical characteristics, such as speed and auto-negotiation is configured and initiates operator notifications (traps) if errors exist. Consistency of MC-LAG configuration (system-id, administrative-key and system-priority) is provided. Similarly, load-balancing mode of operation must be consistently configured on both nodes.
 - Traffic over the signaling link is encrypted using a user configurable message digest key.
- MC-LAG function provides active/stand-by status to other software applications in order to build a reliable solution.

Figure 41 and Figure 42 show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
 - both end-points attached with MC-LAG
 - one end-point attached
- Dual-homing to local PE pair
 - both end-points attached with MC-LAG
 - one end-point attached with MC-LAG

– both end-points attached with MC-LAG to two overlapping pairs

Figure 41 MC-LAG Layer 2 Dual Homing to Remote PE Pairs




Fig_6

Figure 42 MC-LAG Layer 2 Dual Homing to Local PE-Pairs




Fig_7

The forwarding behavior of the nodes abide by the following principles. Note that logical destination (actual forwarding decision) is primarily determined by the service (VPLS or VLL) and the principle below applies only if destination or source is based on MC-LAG:

- Packets received from the network will be forwarded to all local active links of the given destination-sap based on conversation hashing. In case there are no local active links, the packets will be cross-connected to inter-chassis pseudowire.

- Packets received from the MC-LAG sap will be forwarded to active destination pseudo-wire or active local links of destination-sap. In case there are no such objects available at the local node, the packets will be cross-connected to inter-chassis pseudowire.

2.7.8.2 MC-LAG and Subscriber Routed Redundancy Protocol (SRRP)


MC-LAG and SRRP enable dual-homed links from any IEEE 802.1ax (formerly 802.3ad) standards-based access device (for example, a IP DSLAM, Ethernet switch or a Video on Demand server) to multiple Layer 2/3 or Layer 3 aggregation nodes. In contrast with slow recovery mechanisms such as Spanning Tree, multi-chassis LAG provides synchronized and stateful redundancy for VPN services or triple play subscribers in the event of the access link or aggregation node failing, with zero impact to end users and their services.

Refer to the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for information about SRRP.

2.7.8.3 Point-to-Point (p2p) Redundant Connection Across Layer 2/3 VPN Network

[Figure 43](#) shows the connection between two multi-service access nodes (MSANs) across a network based on Layer 2/3 VPN pseudo-wires. The connection between MSAN and a pair of PE routers is realized by MC-LAG. From an MSAN perspective, a redundant pair of PE routers acts as a single partner in LACP negotiation. At any time, only one of the routers has an active link in a specified LAG. The status of LAG links is reflected in status signaling of pseudo-wires set between all participating PEs. The combination of active and stand-by states across LAG links as well as pseudo-wires gives only one unique path between a pair of MSANs.

Figure 43 P2P Redundant Connection Through a Layer 2 VPN Network



OSSG116


Note that the configuration in [Figure 43](#) shows one particular configuration of VLL connections based on MC-LAG, particularly the VLL connection where two ends (SAPs) are on two different redundant-pairs. In addition to this, other configurations are possible, such as:

- Both ends of the same VLL connections are local to the same redundant-pair.
- One end VLL endpoint is on a redundant-pair the other on single (local or remote) node.

2.7.8.4 DSLAM Dual Homing in Layer 2/3 TPSDA Model

Figure 44 shows a network configuration where DSLAM is dual homed to pair of redundant PEs by using MC-LAG. Inside the aggregation network redundant-pair of PEs is connecting to VPLS service which provides reliable connection to single or pair of Broadband Service Routers (BSRs).

Figure 44 DSLAM Dual-Homing Using MC-LAG



OSSG110

MC-LAG and pseudo-wire connectivity, PE-A and PE-B implement enhanced subscriber management features based on DHCP-snooping and creating dynamic states for every subscriber-host. As in any point of time there is only one PE active, it is necessary to provide the mechanism for synchronizing subscriber-host state-information between active PE (where the state is learned) and stand-by PE. In addition, VPLS core must be aware of active PE in order to forward all subscriber traffic to a PE with an active LAG link. The mechanism for this synchronization is outside of the scope of this document.

2.8 G.8031 Protected Ethernet Tunnels

The Nokia PBB implementation offers the capability to use core Ethernet tunnels compliant with ITU-T G.8031 specification to achieve 50 ms resiliency for failures in a native Ethernet backbone. For further information regarding Ethernet tunnels, refer to “G.8031 Protected Ethernet Tunnels” in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide*.

2.9 G.8032 Protected Ethernet Rings

Ethernet ring protection switching offers ITU-T G.8032 specification compliance to achieve resiliency for Ethernet Layer 2 networks. Similar to G.8031 linear protection (also called Automatic Protection Switching (APS)), G.8032 (Eth-ring) is also built on Ethernet OAM and often referred to as Ring Automatic Protection Switching (R-APS).

For further information regarding Ethernet rings, Refer to “G.8032 Protected Ethernet Rings” in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Services Overview Guide*.

2.10 Ethernet Port Monitoring

Ethernet ports can record and recognize various medium statistics and errors. There are two main types of errors:

- **Frame Based** — Frame based errors are counted when the arriving frame has an error that means the frame is invalid. These types of errors are only detectable when frames are presents on the wire.
- **Symbol Based** — Symbol errors are invalidly encoded symbols on the physical medium. Symbols are always present on an active Ethernet port regardless of the presence of frames.

CRC-Monitor and Symbol-Monitor allows the operator to monitor ingress error conditions on the Ethernet medium and compare these error counts to the thresholds. CRC-Monitor monitors CRC errors. Symbol-Monitor monitors symbol errors. Symbol Error is not supported on all Ethernet ports. Crossing a signal degrade (SD) threshold will cause a log event to be raised. Crossing the configured signal failure (SF) threshold will cause the port to enter an operation state of down. The operator may consider the configuration of other protocols to convey the failure, through timeout conditions.

The error rates are in the form of $M \cdot 10^E - N$. The operator has the ability to configure both the threshold (N) and a multiplier (M). By default if the multiplier is not configured the multiplier is 1. As an example, sd-threshold 3 would result in a signal degrade error rate of $1 \cdot 10^E - 3$ (one error per 1000). Changing the configuration to would sd-threshold 3 multiplier 5 result in a signal degrade rate of $5 \cdot 10^E - 3$ (5 errors per 1000). The signal degrade value must be a lower error rate than the signal failure threshold. This threshold can be used to provide notification that the port is operating in a degraded but not failed condition. These do not equate to a bit error rate (BER). CRC-Monitor provides a CRC error rate. Symbol-Monitor provides a symbol error rate.

The configured error thresholds are compared to the operator specified sliding window to determine if one or both of the thresholds have been crossed. Statistics are gathered every second. This means that every second the oldest statistics are dropped from the calculation. The default 10 second sliding window means that at the 11th second, the oldest 1-second statistical data is dropped and the 11th second is included.

Symbol error crossing differs slightly from CRC-based error crossing. The error threshold crossing is calculated based on the window size and the fixed number of symbols that arrive (ingress) on that port during that window. The following configuration demonstrates this concept.

```
config>port>ethernet# info detail
```

```

-----
symbol-monitor
sd-threshold 5 multiplier 5
sf-threshold 3 multiplier 5
no shutdown
exit

show port 2/1/2 ethernet
=====
Ethernet Interface
=====
Description          : 2/1/2
Interface             : 2/1/2
Link-level            : Ethernet
Admin State           : down
Oper State            : down
Physical Link         : No
Single Fiber Mode     : No
IfIndex               : 69271552
Last State Change     : 06/29/2014 05:04:12
Last Cleared Time      : N/A
Phys State Chng Cnt   : 0

Oper Speed            : N/A
Config Speed          : 1 Gbps
Oper Duplex           : N/A
Config Duplex         : full
MTU                   : 9212
Min Frame Length      : 64 Bytes
Hold time up          : 0 seconds
Hold time down        : 0 seconds
DDM Events            : Enabled

Configured Mode       : network
Dot1Q Ethertype       : 0x8100
PBB Ethertype         : 0x88e7
Ing. Pool % Rate      : 100
Ing. Pool Policy      : n/a
Egr. Pool Policy      : n/a
Net. Egr. Queue Pol   : default
Egr. Sched. Pol       : n/a
Auto-negotiate        : true
Oper Phy-tx-clock     : not-applicable
Accounting Policy     : None
Acct Plcy Eth Phys    : None
Egress Rate           : Default
Load-balance-algo     : Default

Encap Type            : null
QinQ Ethertype        : 0x8100
Egr. Pool % Rate      : 100

MDI/MDX               : unknown
Collect-stats         : Disabled
Collect Eth Phys      : Disabled
Ingress Rate          : Default
LACP Tunnel           : Disabled

Down-when-looped      : Disabled
Loop Detected         : False
Use Broadcast Addr    : False

Keep-alive            : 10
Retry                 : 120

Sync. Status Msg.     : Disabled
Tx DUS/DNU            : Disabled
SSM Code Type         : sdh
Rx Quality Level      : N/A
Tx Quality Level      : N/A

Down On Int. Error    : Disabled

CRC Mon SD Thresh     : Disabled
CRC Mon SF Thresh     : Disabled
CRC Mon Window        : 10 seconds

Sym Mon SD Thresh     : 5*10E-5
Sym Mon SF Thresh     : 5*10E-3
Sym Mon Window        : 10 seconds
Tot Sym Mon Errs      : 0

EFM OAM               : Disabled
EFM OAM Link Mon      : Disabled

Configured Address    : 8c:90:d3:a0:c7:42
Hardware Address      : 8c:90:d3:a0:c7:42

```

Transceiver Data

Transceiver Status : not-equipped

Traffic Statistics

	Input	Output
Octets	0	0
Packets	0	0
Errors	0	0
Utilization (300 seconds)	0.00%	0.00%

Port Statistics

	Input	Output
Unicast Packets	0	0
Multicast Packets	0	0
Broadcast Packets	0	0
Discards	0	0
Unknown Proto Discards	0	0

Ethernet-like Medium Statistics

Alignment Errors :	0	Sngl Collisions :	0
FCS Errors :	0	Mult Collisions :	0
SQE Test Errors :	0	Late Collisions :	0
CSE :	0	Excess Collisns :	0
Too long Frames :	0	Int MAC Tx Errs :	0
Symbol Errors :	0	Int MAC Rx Errs :	0
In Pause Frames :	0	Out Pause Frames :	0

The above configuration results in an SD threshold of 5×10^{-5} (0.00005) and an SF threshold of 5×10^{-3} (0.005) over the default 10-second window. If this port is a 1GbE port supporting symbol monitoring then the error rate is compared against 1,250,000,000 symbols (10 seconds worth of symbols on a 1GbE port 125,000,000). If the error count in the current 10 second sliding window is less than 62,500 then the error rate is below the signal degrade threshold and no action is taken. If the error count is between 62,501 and 6,250,000 then the error rate is above signal degrade but has not breached the signal failure signal threshold and a log event will be raised. If the error count is above 6,250,000 the signal failure threshold is crossed and the port will enter an operation state of down. Consider that this is a very simple example meant to demonstrate the function and not meant to be used as a guide for configuring the various thresholds and window times.

A port is not returned to service automatically when a port enters the failed condition as a result of crossing a signal failure threshold for both CRC-Monitor and Symbol-Monitor. Since the port is operationally down without a physical link error monitoring stops. The operator may enable the port using the **shutdown** and **no shutdown port** commands. Other port transition functions like clearing the MDA or slot, removing the cable, and other physical link transition functions.

2.11 802.3ah OAM

802.3ah Clause 57 (**efm-oam**) defines the Operations, Administration, and Maintenance (OAM) sub-layer, which provides mechanisms useful for monitoring link operation such as remote fault indication and remote loopback control. In general, OAM provides network operators the ability to monitor the health of the network and quickly determine the location of failing links or fault conditions. **efm-oam** described in this clause provides data link layer mechanisms that complement applications that may reside in higher layers.

OAM information is conveyed in slow protocol frames called OAM protocol data units (OAMPDUs). OAMPDUs contain the appropriate control and status information used to monitor, test and troubleshoot OAM-enabled links. OAMPDUs traverse a single link, being passed between peer OAM entities, and as such, are not forwarded by MAC clients (like bridges or switches).

The following **efm-oam** functions are supported:

- **efm-oam** capability discovery
- Active and passive modes
- Remote failure indication — Handling of critical link events (link fault, dying gasp, and so on)
- Loopback — A mechanism is provided to support a data link layer frame-level loopback mode. Both remote and local loopback modes are supported
- **efm-oam** PDU tunneling
- High resolution timer for **efm-oam** in 100ms interval (minimum)
- **efm-oam** link monitoring
- Non-zero Vendor Specific Information Field — The 32-bit field is encoded using the format 00:PP:CC:CC and references TIMETRA-CHASSIS-MIB.
 - 00 — Must be zeros
 - PP — Platform type based on the installed IOM from `tmnxHwEquippedPlatform`. Mixed mode deployments may yield different platform values in the same chassis. Since this is IOM-specific, the IOM's unique hardware ID (`tmnxCardHwIndex`) must be included to retrieve the proper value.
 - CC:CC — Chassis type index value from `tmnxChassisType` which is indexed in `tmnxChassisTypeTable`. The table identifies the specific chassis backplane.

The value 00:00:00:00 is sent for all releases that do not support the non-zero value or are unable to identify the required elements. There is no decoding of the peer or local vendor information fields on the network element. The hexadecimal value is included in the **show port port-id ethernet efm-oam** output.

When the **efm-oam** protocol fails to negotiate a peer session or encounters a protocol failure following an established session the *Port State* will enter the *Link Up* condition. This port state is used by many protocols to indicate the port is administratively UP and there is physical connectivity but a protocol, such as **efm-oam**, has caused the ports operational state to enter a DOWN state. A reason code has been added to help discern if the **efm-oam** protocol is the underlying reason for the Link Up condition.

```
show port
=====
Ports on Slot 1
=====
```

Port Id	Admin State	Link State	Port State	Cfg MTU	Oper MTU	LAG/ Bndl Mode	Port Mode	Port Encp	Port Type	C/QS/S/XFP/ MDIMDX
1/1/1	Down	No	Down	1578	1578	-	netw	null	xcme	
1/1/2	Down	No	Down	1578	1578	-	netw	null	xcme	
1/1/3	Up	Yes	Link Up	1522	1522	-	accs	qinq	xcme	
1/1/4	Down	No	Down	1578	1578	-	netw	null	xcme	
1/1/5	Down	No	Down	1578	1578	-	netw	null	xcme	
1/1/6	Down	No	Down	1578	1578	-	netw	null	xcme	

```

# show port 1/1/3
=====
Ethernet Interface
=====
Description          : 10/100/Gig Ethernet SFP
Interface            : 1/1/3
Link-level           : Ethernet
Admin State          : up
Oper State           : down
Reason Down          : efmOamDown
Physical Link        : Yes
Single Fiber Mode    : No
IfIndex              : 35749888
Last State Change    : 12/18/2012 15:58:29
Last Cleared Time    : N/A
Phys State Chng Cnt  : 1

Oper Speed           : N/A
Config Speed         : 1 Gbps
Oper Duplex          : N/A
Config Duplex        : full

MTU                  : 1522
Min Frame Length     : 64 Bytes
Hold time up         : 0 seconds
Hold time down       : 0 seconds
DDM Events           : Enabled

Configured Mode      : access
Dot1Q Ethertype      : 0x8100
PBB Ethertype        : 0x88e7
Ing. Pool % Rate     : 100
Ing. Pool Policy     : n/a
Egr. Pool Policy     : n/a
Net. Egr. Queue Pol  : default
Egr. Sched. Pol      : n/a
Auto-negotiate       : true
Oper Phy-tx-clock    : not-applicable

Encap Type           : QinQ
Qinq Ethertype       : 0x8100
Egr. Pool % Rate     : 100
MDI/MDX              : unknown

```

Accounting Policy	: None	Collect-stats	: Disabled
Acct Plcy Eth Phys	: None	Collect Eth Phys	: Disabled
Egress Rate	: Default	Ingress Rate	: Default
Load-balance-algo	: Default	LACP Tunnel	: Disabled
Down-when-looped	: Disabled	Keep-alive	: 10
Loop Detected	: False	Retry	: 120
Use Broadcast Addr	: False		
Sync. Status Msg.	: Disabled	Rx Quality Level	: N/A
Tx DUS/DNU	: Disabled	Tx Quality Level	: N/A
SSM Code Type	: sdh		
Down On Int. Error	: Disabled		
CRC Mon SD Thresh	: Disabled	CRC Mon Window	: 10 seconds
CRC Mon SF Thresh	: Disabled		
Configured Address	: d8:ef:01:01:00:03		
Hardware Address	: d8:ef:01:01:00:03		

The operator also has the opportunity to decouple the **efm-oam** protocol from the port state and operational state. In cases where an operator wants to remove the protocol, monitor the protocol only, migrate, or make changes the **ignore-efm-state** can be configured in the **port>ethernet>efm-oam** context. When the **ignore-efm-state** command is configured on a port the protocol continues as normal. However, any failure in the protocol state machine (discovery, configuration, time-out, loops, and so on) will not impact the port on which the protocol is active and the optional ignore command is configured. There will only be a protocol warning message if there are issues with the protocol. The default behavior when this optional command is not configured means the port state will be affected by any **efm-oam** protocol fault or clear conditions. Adding and removing this optional ignore command will immediately represent the *Port State* and *Oper State* based on the active configuration. For example, if the **ignore-efm-state** is configured on a port that is exhibiting a protocol error that protocol error does not affect the port state or operational state and there is no *Reason Down* code. If the **ignore-efm-state** is removed from a port with an existing **efm-oam** protocol error, the port will transition to *Link UP*, *Oper Down* with the reason code *efmOamDown*.

2.11.1 OAM Events

The Information OAMPDU is transmitted by each peer at the configured intervals. This OAMPDU performs keepalive and critical notification functions. Various local conditions are conveyed through the setting of the Flags field. The following Critical Link Event defined in IEEE 802.3 Section 57.2.10.1 are supported;

- Link Fault: The PHY has determined a fault has occurred in the receive direction of the local DTE
- Dying Gasp: An unrecoverable local failure condition has occurred
- Critical Event: An unspecified critical event has occurred

The local node can set or unset the various Flag fields based on the operational state of the port, shutdown or activation of the efm-oam protocol or locally raised events. These Flag fields maintain the setting for the continuance of a particular event. Changing port conditions, protocol state or operator intervention may impact the setting of these fields in the Information OAMPDU.

A peer processing the Information OAMPDU can take a configured action when one or more of these Flag fields are set. By default, receiving a set value for any of the Flag fields will cause the local port to enter the previously mentioned *Link Up* port state and an event will be logged. If this default behavior is not desired, the operator may choose to log the event without affecting the local port. This is configurable per Flag field using the options under **config>port>ethernet>efm-oam>peer-rdi-rx**.

2.11.1.1 Link Monitoring

The efm-oam protocol provides the ability to monitor the link for error conditions that may indicate the link is starting to degrade or has reached an error rate that exceeds an acceptable threshold.

Link monitoring can be enabled for three types of frame errors; **errored-frame**, **errored-frame-period** and **errored-frame-seconds**. The **errored-frame** monitor is the number of frame errors compared to the threshold over a window of time. The **errored-frame-period** monitor is the number of frame errors compared to the threshold over a window of number of received packets. This window is checked once per second to see if the window parameter has been reached. The **errored-frame-seconds** monitor is the number of errored seconds compared to the threshold over a window of time. An errored second is any second with a single frame error.

An errored frame is counted when any frame is in error as determined by the Ethernet physical layer, including jabbers, fragments, FCS or CRC and runts. This excludes jumbo frames with a byte count higher than 9212, or any frame that is dropped by the phy layer prior to reaching the monitoring function.

Each frame error monitor functions independently of other monitors. Each of monitor configuration includes an optional signal degrade threshold **sd-threshold**, a signal failure threshold **sf-threshold**, a **window** and the ability to communicate failure events to the peer by setting a Flag field in the Information OAMPDU or the generation of the Event Notification OAMPDU, **event-notification**. The parameters are uniquely configurable for each monitor.

A degraded condition is raised when the configured signal degrade **sd-threshold** is reached. This provides a first level log only action indicating a link could become unstable. This event does not affect the port state. The critical failure condition is raised when the configured **sf-threshold** is reached. By default, reaching the signal failure threshold will cause the port to enter the *Link Up* condition unless the local signal failure **local-sf-action** has been modified to a **log-only** action. Signal degrade conditions for a monitor in signal failed state will be suppressed until the signal failure has been cleared.

The initial configuration or the modification of either of the threshold values will take effect in the current window. When a threshold value for a monitor is modified, all active local events for that specific monitor will be cleared. The modification of the threshold acts the same as the **clear** command described later in this section.

Notification to the peer is required to ensure the action taken by the local port detecting the error and its peer are synchronized. If peers do not take the same action then one port may remain fully operational while the other enters a non-operational state. These threshold crossing events do not shutdown the physical link or cause the protocol to enter a non-operational state. The protocol and network element configuration is required to ensure these asymmetrical states do not occur. There are two options for exchanging link and event information between peers; Information OAMPDU and the Event Notification OAMPDU.

As discussed earlier, the Information OAMPDU conveys link information using the Flags field; dying gasp, critical link and link fault. This method of communication has a number of significant advantages over the Event Notification OAMPDU. The Information OAMPDU is sent at every configured **transmit-interval**. This will allow the most recent information to be sent between peers, a critical requirement to avoid asymmetrical forwarding conditions. A second major advantage is interoperability with devices that do not support Link Monitoring and vendor interoperability. This is the lowest common denominator that offers a robust communication to convey link event information. Since the Information OAMPDU is already being sent to maintain the peering relationship this method of communication adds no additional overhead. The **local-sf-action** options allow the dying gasp and critical event flags to be set in the Information OAMPDU when a signal failure threshold is reached. It is suggested that this be used in place of or in conjunction with Event Notification OAMPDU.

Event Notification OAMPDU provides a method to convey very specific information to a peer about various Link Events using Link Event TLVs. A unique Event Notification OAMPDU will be generated for each unique frame error event. The intention is to provide the peer with the Sequence Number, Event Type, Timestamp, and the local information that caused the generation of the OAMPDU; window, threshold, errors and error running total and event running total specific to the port.

- Sequence Number: The unique identification indicating a new event.

- Window: The size of the unique measurement period for the error type. The window is only checked at the end. There is not mid-window checking.
- Threshold: The value of the configured sf-threshold
- Errors: The errors counted in that specific window
- Error Running Total: The number of errors accumulated for that event type since monitoring started and the protocol and port have been operational or a reset function has occurred
- Event Running Total: The number of events accumulated for that event type since the monitoring started and the protocol and port have been operational

By default, the Event Notification OAMPDU is generated by the network element detecting the signal failure event. The Event Notification OAMPDU is sent only when the initial frame event occurs. No Event Notification OAMPDU is sent when the condition clears. A port that has been operationally affected as a result of a Link Monitoring frame error event must be recovered manually. The typical recovery method is to shutdown the port and no shutdown the port. This clears all events on the port. Any function that affects the port state, physical fiber pull, soft or hard reset functions, protocol restarts, and so on, also clears all local and remote events on the affected node experiencing the operation. None of these frame errors recovery actions will cause the generation of the Event Notification OAMPDU. If the chosen recovery action is not otherwise recognized by the peer and the Information OAMPDU Flag fields have not been configured to maintain the current event state, there is a high probability that the ports will have different forwarding states, notwithstanding any higher level protocol verification that may be in place.

A burst of between one and five Event Notification OAMPDU packets may be sent. By default, only a single Event Notification OAMPDU is generated, but this value can be changed under the **local-sf-action** context. An Event Notification OAMPDU will only be processed if the peer had previously advertised the EV capability. The EV capability is an indication the remote peer supports link monitoring and may send the Event Notification OAMPDU.

The network element receiving the Event Notification OAMPDU will use the values contained in the Link event TLVs to determine if the remote node has exceeded the failure threshold. The locally configured action will determine how and if the local port is affected. By default, processing of the Event Notification OAMPDU is log only and does not affect the port state. By default, processing of the Information OAMPDU Flag fields is port affecting. When Event Notification OAMPDU has been configured as port affecting on the receiving node, action is only taken when errors are equal to or above the threshold and the threshold value is not zero. No action is taken when the errors value is less than the threshold or the threshold is zero.

Symbol error, **errored-symbols**, monitoring is also supported but requires specific hardware revisions and the appropriate code release. The symbol monitor differs from the frame error monitors. Symbols represent a constant load on the Ethernet wire whether service frames are present or not. This means the optional signal degrade threshold **sd-threshold** has an additional purpose when configured as part of the symbol error monitor. When the signal degrade threshold is not configured, the symbol monitor acts similar to the frame error monitors, requiring manual intervention to clear a port that has been operationally affected by the monitor. When the optional signal degrade threshold is configured, it again represents the first level warning. However, it has an additional function as part of the symbol monitor. If a signal failure event has been raised, the configured signal degrade threshold becomes the equivalent to a lowering threshold. If a subsequent window does not reach the configured signal degrade threshold then the previous event will be cleared and the previously affected port will be returned to service without operator intervention. This return to service will automatically clear any previously set Information OAMPDU Flags fields set as a result of the signal failure threshold. The Event Notification OAMPDU will be generated with the symbol error Link TLV that contains an error count less than the threshold. This will indicate to the peer that initial problem has been resolved and the port should be returned to service.

The **errored-symbol** window is a measure of time that is automatically converted into the number of symbols for that specific medium for that period of time. The standard MIB entries “dot3OamErrSymPeriodWindowHi” and “dot3OamErrSymPeriodWindowLo” are marked as read-only instead of read-write. These values cannot be configured directly. The configuration of the **window** converts the time and programs the two MIB values in an appropriate manner. Both the configured **window** and the number of symbols will be displayed under the **show port port-id ethernet efm-oam** command.

```
show port 1/1/1 ethernet efm-oam
=====
Ethernet Oam (802.3ah)
=====
Admin State      : up
Oper State       : operational
Mode             : active
Pdu Size         : 1518
Config Revision  : 0
Function Support  : LB
Transmit Interval : 1000 ms
Multiplier       : 5
Hold Time        : 0
Tunneling        : false
Loop Detected    : false
Grace Tx Enable   : true (inactive)
Grace Vendor OUI  : 00:16:4d
Dying Gasp on Reset : true (inactive)
Soft Reset Tx Act : none
Trigger Fault     : none
Vendor OUI        : 00:16:4d (alu)
Vendor Info       : 00:01:00:02
```

```

Peer Mac Address      : d8:1c:01:02:00:01
Peer Vendor OUI       : 00:16:4d (alu)
Peer Vendor Info      : 00:01:00:02
Peer Mode             : active
Peer Pdu Size         : 1518
Peer Cfg Revision     : 0
Peer Support          : LB
Peer Grace Rx         : false
Loopback State        : None
Loopback Ignore Rx    : Ignore
Ignore Efm State      : false
Link Monitoring       : disabled
Peer RDI Rx
  Critical Event      : out-of-service
  Dying Gasp          : out-of-service
  Link Fault          : out-of-service
  Event Notify        : log-only
Local SF Action
  Event Burst         : 1
  Port Action         : out-of-service
  Dying Gasp          : disabled
  Critical Event      : disabled
Discovery
  Ad Link Mon Cap    : yes

Errored Frame
  Enabled             : no
  Event Notify        : enabled
  SF Threshold        : 1
  SD Threshold        : disabled (0)
  Window              : 10 ds
Errored Frame Period
  Enabled             : no
  Event Notify        : enabled
  SF Threshold        : 1
  SD Threshold        : disabled (0)
  Window              : 1488095 frames
Errored Symbol Period
  Enabled             : no
  Event Notify        : enabled
  SF Threshold        : 1
  SD Threshold        : disabled (0)
  Window (time)       : 10 ds
  Window (symbols)    : 125000000
Errored Frame Seconds Summary
  Enabled             : no
  Event Notify        : enabled
  SF Threshold        : 1
  SD Threshold        : disabled (0)
  Window              : 600 ds

=====
Active Failure Ethernet OAM Event Logs
=====
Number of Logs : 0
=====

Ethernet Oam Statistics
=====

```

	Input	Output
Information	238522	238522
Loopback Control	0	0
Unique Event Notify	0	0
Duplicate Event Notify	0	0
Unsupported Codes	0	0
Frames Lost		0

```

=====

```

A **clear** command “**clear port *port-id* ethernet efm-oam events [local | remote]**” has been added to clear port affecting events on the local node on which the command is issued. When the optional [**local | remote**] options are omitted, both local and remote events will be cleared for the specified port. This command is not specific to the link monitors as it clears all active events. When local events are cleared, all previously set Information OAMPDU Flag fields will be cleared regardless of the cause of the event that set the Flag field.

In the case of symbol errors only, if Event Notification OAMPDU is enabled for symbol errors and a local symbol error signal failure event exists at the time of the clear, the Event Notification OAMPDU will be generate with an error count of zero and a threshold value reflecting the local signal failure threshold. An error value lower than the threshold value indicates the local node is not in a signal failed state. The Event Notification OAMPDU is not generated in the case where the **clear** command clears local frame error events. This is because frame error event monitors will only act on an Event Notification OAMPDU when the error value is higher than the threshold value, a lower value is ignored. As stated previously, there is no automatic return to service for frame errors.

If the clear command clears remote events, events conveyed to the local node by the peer, no notification is generated to the peer to indicate a clear function has been performed. Since the Event Notification OAMPDU is only sent when the initial event was raised, there is no further Event Notification and blackholes can result. If the Information OAMPDU Flag fields are used to ensure a constant refresh of information, the remote error will be reinstated as soon as the next Information OAMPDU arrives with the appropriate Flag field set.

Local and remote efm-oam port events are stored in the efm-oam event logs. These logs maintain and display active and cleared signal failure degrade events. These events are interacting with the efm-oam protocol. This logging is different than the time stamped events for information logging purposes included with the system log. To view these events, the **event-log** option has been added to the **show port *port-id* ethernet efm-oam** command. This includes the location, the event type, the counter information or the decoded Network Event TLV information, and if the port has been affected by this active event. A maximum of 12 port events will be retained. The first three indexes are reserved for the three Information Flag fields, dying gasp, critical link, and link fault. The other nine indexes will maintain the current state for the various error monitors in a most recent behavior and events can wrap the indexes, dropping the oldest event.

In mixed environments where Link Monitoring is supported on one peer but not the other the following behavior is normal, assuming the Information OAMPDU has been enabled to convey the monitor fault event. The arriving Flag field fault will trigger the efm-oam protocol on the receiving unsupportive node to move from operational to “send local and remote”. The protocol on the supportive node that set the Flag field to convey the fault will enter the “send local and remote ok” state. The supportive

node will maintain the Flag field setting until the condition has cleared. The protocol will recover to the operational state once the original event has cleared; assuming no other fault on the port is preventing the negotiation from progressing. If both nodes were supportive of the Link Monitoring process, the protocol would remained operational.

In summary, Link monitors can be configured for frame and symbol monitors (specific hardware only). By default, Link Monitoring and all monitors are shutdown. When the Link Monitoring function is enabled, the capability (EV) will be advertised. When a monitor is enabled, a default window size and a default signal failure threshold are activated. The local action for a signal failure threshold event is to shutdown the local port. Notification will be sent to the peer using the Event Notification OAMPDU. By default, the remote peer will not take any port action for the Event Notification OAMPDU. The reception will only be logged. It is suggested the operator evaluate the various defaults and configure the **local-sf-action** to set one of the Flag fields in the Information OAMPDU using the **info-notifications** command options when fault notification to a peer is required. Non-Nokia vendor specific information will not be processed.

2.11.1.1 Capability Advertising

A supported capability, sometimes requiring activation, will be advertised to the peer. The EV capability is advertisement when Link Monitoring is active on the port. This can be disabled using the optional command **no link-monitoring** under the **config>port>ethernet>efm-oam>discovery>advertise-capabilities**.

2.11.2 Remote Loopback

EFM OAM provides a link-layer frame loopback mode that can be remotely controlled.

To initiate remote loopback, the local EFM OAM client sends a loopback control OAM PDU by enabling the OAM remote-loopback command. After receiving the loopback control OAM PDU, the remote OAM client puts the remote port into local loopback mode.

To exit remote loopback, the local EFM OAM client sends a loopback control OAM PDU by disabling the OAM remote-loopback command. After receiving the loopback control OAM PDU, the remote OAM client puts the port back into normal forwarding mode.

During remote loopback test operation, all frames except EFM OAM PDUs are dropped at the local port for the receive direction, where remote loopback is enabled. If local loopback is enabled, then all frames except EFM OAM PDUs are dropped at the local port for both the receive and transmit directions. This behavior may result in many protocols (such as STP or LAG) resetting their state machines.

When a port is in loopback mode, service mirroring will not work if the port is a mirror-source or a mirror-destination.

2.11.3 802.3ah OAM PDU Tunneling for Epipe Service

Nokia routers support 802.3ah. Customers who subscribe to Epipe service treat the Epipe as a wire, so they demand the ability to run 802.3ah between their devices which are located at each end of the Epipe.

This feature only applies to port-based Epipe SAPs because 802.3ah runs at port level not VLAN level. Hence, such ports must be configured as null encapsulated SAPs.

When OAM PDU tunneling is enabled, 802.3ah OAM PDUs received at one end of an Epipe are forwarded through the Epipe. 802.3ah can run between devices that are located at each end of the Epipe. When OAM PDU tunneling is disabled (by default), OAM PDUs are dropped or processed locally according to the **efm-oam** configuration (**shutdown** or **no shutdown**).

Enabling 802.3ah for a specific port and enabling OAM PDU tunneling for the same port are mutually exclusive. Enforcement is performed at the CLI level.

2.11.3.1 802.3ah Grace Announcement

Support for vendor-specific soft reset graceful recovery has been added to the SR OS implementation of the EFM-OAM protocol. This is configured using the **grace-tx-enable** command under the **config>system>ethernet>efm-oam** and the **config>port>ethernet>efm-oam** contexts. This feature is not enabled by default. When this functionality is enabled, the EFM-OAM protocol does not enter a non-operational state when both nodes acknowledge the grace function. The ports associated with the hardware that has successfully executed the soft reset will clear all local and remote events. The peer that acknowledges the graceful restart

procedure for EFM-OAM will clear all remote events that it received from the peer that performed the soft reset. The local events will not be cleared on the peer that has not undergone soft reset. The Information OAM PDU Flag fields are critical in propagating the local event to the peer. The Event Notification OAM PDU will not be sent because it is only sent when the event is initially raised.

A vendor-specific Grace TLV will be included in the Information PDU generated as part of the 802.3ah OAM protocol when a network element undergoes an ISSU function. Nodes that support the Soft Reset messaging functions will allow the local node to generate the grace TLV.

The grace TLV informs a remote peer that the negotiated interval and multiplier should be ignored and the new 900s timeout interval should be used to timeout the session. The peer receiving the Grace TLV must be able to parse and process the vendor specific messaging.

The new command **grace-tx-enable** has been introduced to enable this functionality. This command exists at two levels of the hierarchy, system level and port level. By default this functionality is enabled on the port. At the system level this command defaults to disabled. In order to enable this functionality both the port and the system commands must be enabled. If either is not enabled then the combination will not allow those ports to generate the vendor specific Grace TLV. This functionality must be enabled at both the system and port level prior to the ISSU or soft reset function. If this is enabled during a soft reset or after the ISSU function is already in progress it will have no affect during that window. Both Passive and Active 802.3ah OAM peers can generate the Grace TLV as part of the informational PDU.

There is no command to enable this on the receiving node. As long as the receiver understands and can parse the Grace TLV it will enter the grace mode of operation.

The basic protocol flow below helps demonstrate the interaction between passive-active and active-active peer combinations supporting the Grace TLV. In the first diagram the passive node is entering an ISSU on a node that supports soft reset capabilities.

In [Figure 45](#) and [Figure 46](#), the Active node is experiencing the ISSU function on a node that supports soft reset capabilities.

Figure 45 Grace TLV Passive Node with Soft Reset





Figure 46 Grace TLV Active Node with Soft Reset

The difference between the two is subtle but important. When an active node performs this function it will generate an Informational TLV with the Local TLV following the successful soft reset. When it receives the Information PDU with the Grace Ack it will send its own Information PDU with both Local and Remote TLV completed. This will complete the protocol restart. When a passive node is reset the passive port will wait to receive the 802.3ah OAM protocol before sending its own Information PDU with both the Local and Remote TLV thus completing the protocol restart.

The renegotiation process allows the node which experienced the soft reset to rebuild the session without having to restart the session from the discovery phase. This significantly reduces the impact of the native protocol on data forwarding.

Any situation that could cause the renegotiation to fail will force the protocol to revert to the discovery phase and fail the graceful restart. During a Major ISSU when the EFM-OAM session is held operational by the Grace function, if the peer MAC address of the session changes, there will be no log event raised for the MAC address change.

The vendor-specific grace function benefits are realized when both peers support the transmitting, receiving and processing of the vendor-specific Grace TLV. In the case of mixed code versions, products, or vendor environments, a standard EFM-OAM message to the peer can be used to instruct the peer to treat the session as failed. When the command **dying-gasp-tx-on-reset** is active on a port, the soft reset function triggers ETH-OAM to set the dying gasp flag or critical event flag in the Information OAMPDU. An initial burst of three Informational OAM PDUs will be sent using a 1-second spacing, regardless of the protocol interval. The peer may process these flags to affect its port state and take the appropriate action. The control of the local port state where the soft reset is occurring is left to the soft reset function. This EFM-OAM function does not affect local port state. If the peer has acted on the exception flags and affected its port state, then the local node must take an action to inform the upstream nodes that a condition has occurred and forwarding is no longer possible. Routing protocols like ISIS and OSPF overload bits are typically used in routed environments to accomplish this notification.

This feature is similar to **grace-tx-enable**. Intercepting system messaging, when the feature is active on a port (enabled both at the port and at the system level) and when the messaging occurs, is a similar concept. However, because the **dying-gasp-tx-on-reset** command is not a graceful function it is interruptive and service affecting. Using **dying-gasp-tx-on-reset** requires peers to reestablish the peering session from an initial state, not rebuild the state from previous protocol information. The transmission of the dying gasp or the critical event commences when the soft reset occurs and continues for the duration of the soft reset.

If both functions are active on the same port, the **grace-tx-enable** function is preferred if the peer is setting and sending the Vendor OUI to 00:16:4d (ALU) in the Information OAMPDU. In this situation, the dying gasp function will not be invoked. A secondary Vendor OUI can be configured using the **grace-vendor-oui oui** command, should an additional Vendor OUI prefer to support the reception, parsing, and processing of the vendor-specific grace message instead of the dying gasp. If only one of those functions is active on the port then that specific function will be called. The grace function should not be enabled if the peer Vendor OUI is equal to 00:16:4d (ALU) and the peer does not support the grace function.

ETH-OAM allows generation of a fault condition by using the **trigger-fault {dying-gasp | critical-event}** command. This sets the appropriate flag fields in the Information OAMPDU and transitions a previously operational local port to Link Up. Removing this command from the configuration stops the flags from being set and allows the port to return to service, assuming no other faults would prevent this resumption of service. In cases where a port must be administratively shut down, this command can be used to signal a peer using the EFM-OAM protocol, and the session should be considered failed.

These features do not support the clearing of an IOM which does not trigger a soft reset. IOM clearing is a forceful event that does not trigger graceful protocol renegotiation.

A number of show commands have been enhanced to help operators determine the state of the 802.3ah OAM Grace function and whether or not the peer is generating or receiving the Grace TLV.

System level information can be viewed using the **show system info** command.

```
show system information
=====
System Information
=====
System Name       : system-name
System Type       : 7750 SR-12
System Version    : 11.0r4
System Contact    :
System Location   :
System Coordinates:
System Active Slot: A
System Up Time    : 62 days, 20:29:48.96 (hr:min:sec)

...snip...

EFM OAM Grace Tx Enable: False
=====
```

EFM OAM Grace Tx Enable:

- False — The system level functionality is not enabled. Grace will not be generated on any ports regardless of the state of the option on the individual ports
- True — The system level functionality is enabled and the determination of whether to send grace is based on the state of the option configured at the port level

Individual ports also contain information about the current port configuration and whether or not the Grace TLV is being sent or received.

Grace Tx Enable has two enable states with the current state in brackets to the right.

- False — The port level functionality is not enabled. Grace will not be generated on the port regardless of the state of the option at the system level.
- True — The port level functionality is enabled and the determination of whether to send grace is based on the state of the option configured at the system level
 - (inactive) Not currently sending Grace TLV
 - (active) Currently sending the Grace TLV as part of the Information PDU

Peer Grace Rx

- False — Not receiving Grace TLV from the peer
- True — Receiving Grace TLV from the peer

2.12 MTU Configuration Guidelines

Observe the following general rules when planning your service and physical MTU configurations:

- The router must contend with MTU limitations at many service points. The physical (access and network) port, service, and SDP MTU values must be individually defined.
- Identify the ports that will be designated as network ports intended to carry service traffic.
- MTU values should not be modified frequently.
- MTU values must conform to both of the following conditions:
 - The service MTU must be less than or equal to the SDP path MTU.
 - The service MTU must be less than or equal to the access port (SAP) MTU.
- When the network group encryption (NGE) feature is enabled, additional bytes due to NGE packet overhead must be considered. Refer to the “NGE Packet Overhead and MTU Considerations” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information.

2.12.1 Default MTU Values

[Table 36](#) shows the default MTU values which are dependent upon the (sub-) port type, mode, and encapsulation.

Table 36 MTU Default Values

Port Type	Mode	Encap Type	Default (bytes)
Ethernet	access	null	1514
Ethernet	access	dot1q	1518
Fast Ethernet	network	—	1514
Other Ethernet	network	—	9212 ¹
SONET path or TDM channel	access	BCP-null	1518
SONET path or TDM channel	access	BCP-Dot1q	1522
SONET path or TDM channel	access	IPCP	1502

Table 36 MTU Default Values (Continued)

Port Type	Mode	Encap Type	Default (bytes)
SONET path or TDM channel	network	—	9208
SONET path or TDM channel	access	frame-relay	1578
SONET path or TDM channel	access	atm	1524

Note:

1. The default MTU for Ethernet ports other than Fast Ethernet is actually the lesser of 9212 and any MTU limitations imposed by hardware which is typically 16K.

2.12.2 Modifying MTU Defaults

MTU parameters must be modified on the service level as well as the port level.


- The service-level MTU parameters configure the service payload (Maximum Transmission Unit – MTU) in bytes for the service ID overriding the service-type default MTU.
- The port-level MTU parameters configure the maximum payload MTU size for an Ethernet port or SONET/SDH SONET path (sub-port) or TDM port/channel, or a channel that is part of a multilink bundle or LAG.

The default MTU values must be modified to ensure that packets are not dropped due to frame size limitations. The service MTU must be less than or equal to both the SAP port MTU and the SDP path MTU values. When an SDP is configured on a network port using default port MTU values, the operational path MTU can be less than the service MTU. In this case, enter the `show service sdp` command to check the operational state. If the operational state is down, then modify the MTU value accordingly.

2.12.3 Configuration Example

In order for the maximum length service frame to successfully travel from a local ingress SAP to a remote egress SAP, the MTU values configured on the local ingress SAP, the SDP (GRE or MPLS), and the egress SAP must be coordinated to accept the maximum frame size the service can forward. For example, the targeted MTU values to configure for a distributed Epipe service (ALA-A and ALA-B) are shown in [Figure 47](#).

Figure 47 MTU Configuration Example



Since ALA-A uses Dot1q encapsulation, the SAP MTU must be set to 1518 to be able to accept a 1514 byte service frame (see [Table 36](#) for MTU default values). Each SDP MTU must be set to at least 1514 as well. If ALA-A’s network port (2/1/1) is configured as an Ethernet port with a GRE SDP encapsulation type, then the MTU value of network ports 2/1/1 and 3/1/1 must *each* be at least 1556 bytes (1514 MTU + 28 GRE/Martini + 14 Ethernet). Finally, the MTU of ALA-B’s SAP (access port 4/1/1) must be at least 1514, as it uses null encapsulation.

[Table 37](#) shows sample MTU configuration values.

Table 37 MTU Configuration Example Values

	ALA-A		ALA-B	
	Access (SAP)	Network	Network	Access (SAP)
Port (slot/MDA/port)	1/1/1	2/1/12	3/1/1	4/1/1
Mode type	dot1q	network	network	null
MTU	1518	1556	1556	1514

2.13 Deploying Preprovisioned Components

When a card, CMA, MDA, XCM or XMA is installed in a preprovisioned slot, the device detects discrepancies between the preprovisioned card type configurations and the types actually installed. Error messages display if there are inconsistencies and the card will not initialize.

When the proper preprovisioned cards are installed into the appropriate chassis slot, alarm, status, and performance details will display.

2.14 Setting Fabric Speed

The **set-fabric-speed** command sets fabric speed for the 7750 SR-7/12/12e and the 7450 ESS-7/12. This tools command is added with the introduction of the T3 generation of fabric switching.

2.14.1 fabric-speed-a (SFM5/CPM5)

The 7750 SR-7/12/12e and 7450 ESS-7/12 chassis defaults to the **fabric-speed-a** when initially deployed. The chassis operates as: 100 Gb/s per slot for ESS/SR-7/12 and 200 Gb/s per slot for SR-12e. This permits a mixture of FP2/FP3 based cards to co-exist.

2.14.2 fabric-speed-b (SFM5/CPM5)

The **fabric-speed-b** parameter enables the 7750 SR-7/12 and 7450 ESS-7/12 to operate at up to 200 Gb/s, and the 7750 SR-12e to operate up to 400 Gb/s. All cards in the system are required to be T3 based (FP3 IMM and/or IOM3-XP-C or newer). The system does not support any FP2 based cards when the chassis is set to **fabric-speed-b**.



Note: For the 7750 SR-7/12 and 7450 ESS-7/12, the chassis must have a manufacturing date of 2008 or later (**show chassis detail**) to set **fabric-speed-b**.

2.14.3 fabric-speed-c (SFM5/CPM5)

Enables the 7950 XRS-20/20e or 7750 SR-12e to use both FP3- and FP4-based cards. This speed is mandatory if FP4 cards are used.

2.15 Versatile Service Module

2.15.1 VSM-CCA-XP

The VSM-CCA-XP MDA offers a hybrid mode for simplified provisioning and a high capacity VSM when inserted on IOM cards. The complete forwarding path bandwidth (in this case 25 Gb/s) is available, allowing single conversations up to 25 Gb/s on a single MDA.

A VSM-CCA-XP has two ports, port x/x/1 and port x/x/2, which are internally connected. Configuration is very similar to a physical loop-back port using Ethernet hybrid ports with dot1Q encapsulation. The use of the Ethernet VLAN tag connects the SAPs.

The following constraints apply to the configuration of the VSM-CCA-XP:

- While LAG is available LACP is not allowed.
- Ethernet CFM is only available when Eth-Rings are configured on the VSM (Ethernet rings use Ethernet MEPS for control).
- ATM/Ethernet Last-Mile Aware QoS for Broadband Network Gateway is not supported.

The following is a sample configuration for an MDA provisioned as a VSM-CCA-XP MDA.

```
*A:PE-1>config>card# info
-----
      card-type iom3-xp
      mda 1
        mda-type m10-1gb-xp-sfp
        no shutdown
      exit
      mda 2
        mda-type vsm-cca-xp
        no shutdown
      exit
      no shutdown
-----

*A:PE-1>config>card# exit all
*A:PE-1# show mda
=====
MDA Summary
=====
```

Slot	Mda	Provisioned Type Equipped Type (if different)	Admin State	Operational State
1	1	m10-1gb-xp-sfp	up	up
	2	vsm-cca-xp	up	up

```
=====
*A:PE-1#
```

The following is a sample VSM-CCM-XP configuration for ports:

```
port 1/2/1
  ethernet
  exit
  no shutdown
exit
port 1/2/2
  ethernet
  exit
  no shutdown
exit
```

Port and Ethernet QoS parameters may be configured as they are for a physical port. The Ethernet on VSM-CCA-XP has a reduced set of features. For example, dot1Q is the only supported encapsulation. LACP cannot be configured on LAGs using the port.

The ports may be used directly by service SAP in the case of a single loop back. If resiliency is required, or more capacity is needed, a LAG can be configured.

The following is a sample configuration for LAG on a single VSM-CCA-XP MDA:

```
lag 1
  mode hybrid
  encap-type dot1q
  port 1/2/1 // VSM-CCA-XP
  no shutdown
exit
lag 2
  mode hybrid
  encap-type dot1q
  port 1/2/2 // VSM-CCA-XP
  no shutdown
exit
```

The following is a sample for a VPLS service equivalent using the LAG port:

```
vpls 121 customer 1 create
  stp
  shutdown
exit
sap lag-1:1001 create // Connect using VLAN Tag 1001
exit
no shutdown
...
exit
```


The following is a sample for an IES service equivalent to the configuration:

```
ies 122 customer 1 create
  interface "Loopback" create
    address 10.1.1.1/24
    sap lag-2:1001 create
      ingress
        qos 3
      exit
      egress
        qos 1010
      exit
    exit
  exit
  ...
  no shutdown
exit
```

2.16 Configuration Process Overview

Figure 48 displays the process to provision chassis slots, cards, MDAs, and ports.

Figure 48 Slot, Card, MDA, and Port Configuration and Implementation Flow



2.17 Configuration Notes

The following information describes provisioning caveats:

- If a card or MDA type is installed in a slot provisioned for a different type, the card will not initialize.
- A card or MDA installed in an unprovisioned slot remains administratively and operationally down until the card type and MDA is specified.
- Ports cannot be provisioned until the slot, card and MDA type are specified.
- cHDLC does not support HDLC windowing features, nor other HDLC frame types such as S-frames.
- cHDLC operates in the HDLC Asynchronous Balanced Mode (ABM) of operation.
- APS configuration rules:
 - A physical port (either working or protection) must be shutdown before it can be removed from an APS group port.
 - For a single-chassis APS group, a working port must be added first. Then a protection port can be added or removed at any time.
 - A protection port must be shutdown before being removed from an APS group.
 - A path cannot be configured on a port before the port is added to an APS group.
 - A working port cannot be removed from an APS group until the APS port path is removed.
 - When ports are added to an APS group, all path-level configurations are available only on the APS port level and configuration on the physical member ports are blocked.
 - For APS-protected bundles, all members of a working bundle must reside on the working port of an APS group. Similarly all members of a protecting bundle must reside on the protecting circuit of that APS group.

2.18 Configuring Physical Ports with CLI

This section provides information to configure cards, MDAs, and ports.

2.18.1 Preprovisioning Guidelines

SR OSs have a console port, either located on the CPM or CCM, or integrated into the chassis (on the 7750 SR-c4 models), to connect terminals to the router.

Configure parameters from a system console connected to a router console port, using Telnet to access a router remotely or SSH to open a secure shell connection.

2.18.1.1 Predefining Entities

In order to initialize a card, the chassis slot, line card type, and MDA type must match the preprovisioned parameters. In this context, *preprovisioning* means to configure the entity type (such as the card type, MDA type, port, and interface) that is planned for a chassis slot, card, or MDA. Preprovisioned entities can be installed but not enabled or the slots can be configured but remain empty until populated. *Provisioning* means that the preprovisioned entity is installed and enabled.

You can:

- Pre-provision ports and interfaces after the line card and MDA types are specified.
- Install line cards in slots with no preconfiguration parameters specified. Once the card is installed, the card and MDA types must be specified.
- Install a line card in a slot provisioned for a different card type (the card will not initialize). The existing card and MDA configuration must be deleted and replaced with the current information.

2.18.1.2 Preprovisioning a Port

Before a port can be configured, the slot must be preprovisioned with an allowed card type and the MDA must be preprovisioned with an allowed MDA type. Some recommendations to configure a port include:

- Ethernet

- Configure an access port for customer facing traffic on which services are configured.
An encapsulation type may be specified in order to distinguish services on the port or channel. Encapsulation types are not required for network ports.
To configure an Ethernet access port, see [Ethernet Access Port](#).
- SONET/SDH
 - SONET/SDH can be used only when configuring an OC-3, OC-12, OC-48, OC-192, and OC-768 SONET paths on an appropriate MDA.
To configure a SONET path, see [Configuring SONET/SDH Port Parameters](#).
Configure a network port or channel to participate in the service provider transport or infrastructure network.
Accounting policies can only be associated with network ports/channels and Service Access Ports (SAPs). Accounting policies are configured in the **config>log> accounting-policy** context.
To configure an Ethernet network port, see [Ethernet Network Port](#).
- Channelized
 - Channelized ports can only be configured on channel-capable MDAs, CMAs, or TDM satellites such as the channelized DS-3, channelized OC-3-SFP, channelized OC-12-SFP, or channelized Any Service Any Port MDAs or CMAs.

2.18.1.3 Maximizing Bandwidth Use

Once ports are preprovisioned, Link Aggregation Groups (LAGs), multilink-bundles (IMA), or Bundle Protection Groups (for example IMA BPGs), can be configured to increase the bandwidth available between two nodes.

All physical links or channels in a given LAG/bundle combine to form one logical connection. A LAG/bundle also provides redundancy in case one or more links that participate in the LAG/bundle fail. For command syntax for multilink bundles, see [Configuring Multilink PPP Bundles](#). To configure channelized port for TDM, see [Configuring Channelized Ports](#). To configure channelized port for Sonet/SDH high speed channels (ASAP MDAs only), see [Configuring SONET/SDH Port Parameters](#). For command syntax for LAG, see [Configuring LAG Parameters](#).

2.18.2 Basic Configuration

The most basic configuration must specify the following:

- chassis slot
- line card type (must be an allowed card type)
- MCM slot for the 7750 SR-c4 and SR-c12 only (not required for CMA)
- MCM type the 7750 SR-c4 and SR-c12 only (must be an allowed MCM type)
- MDA slot
- MDA (must be an allowed MDA type)
- specific port to configure

The following is an example of card configuration for the 7750 SR:

```
A:PE-1# admin display-config | match "Card Configuration" post-lines 28
echo "Card Configuration"
#-----
    card 2
        card-type imm5-10gb-xfp
        mda 1
            no shutdown
        exit
        no shutdown
    exit
    card 3
        card-type iom4-e
        mda 1
            mda-type me10-10gb-sfp+
            no shutdown
        exit
        mda 2
            mda-type me1-100gb-cfp2
            no shutdown
        exit
        no shutdown
    exit
    card 5
        card-type imm5-10gb-xfp
        mda 1
            no shutdown
        exit
        no shutdown
    exit
#-----
A:PE-1#
```

The following is an example of card configuration on a 7750 SR-c12:

```
A:Dut-C#
A:Dut-C>config>card# info
-----
```

```
card-type iom-xp
mcm 1
    mcm-type mcm-xp
    no shutdown
exit
mcm 3
    mcm-type mcm-xp
    no shutdown
exit
mcm 7
    mcm-type mcm-xp
    no shutdown
exit
mda 1
    mda-type m4-choc3-as-sfp
    no shutdown
exit
mda 3
    mda-type
    no shutdown
exit
mda 5
    mda-type c8-10/100eth-tx
    no shutdown
exit
mda 7
    mda-type isa-aa
    no shutdown
exit
mda 12
    mda-type c8-10/100eth-tx
    no shutdown
exit
no shutdown
```

A:Dut-C>config>card#

The following is an example of card configurations for the 7950 XRS:

A:7950 XRS-20# configure card 1
A:7950 XRS-20>config>card# info

```
-----
card-type xcm-x20
mda 1
    mda-type cx20-10g-sfp
    no shutdown
exit
mda 2
    mda-type cx2-100g-cfp
    no shutdown
exit
no shutdown
-----
```

2.18.3 Common Configuration Tasks

The following sections are basic system tasks that must be performed.

2.18.3.1 Configuring Cards and MDAs

Card configurations include a chassis slot designation. A slot must be preconfigured with the type of cards and MDAs which are allowed to be provisioned.

The following example shows card and MDA configurations for the 7750 SR or 7450 ESS:

```
#-----
\echo "Card Configuration"
#-----
card 1
  card-type iom3-xp-c
  mda 1
    mda-type m10-1gb+1-10gb
    no shutdown
  exit
  mda 2
    mda-type m1-choc12-as-sfp
    no shutdown
  exit
  no shutdown
exit
```

The following example shows card configurations for the 7950 XRS:

```
A:7950 XRS-20# configure card 1
A:7950 XRS-20>config>card# info
-----
card-type xcm-x20
mda 1
  mda-type cx20-10g-sfp
  no shutdown
exit
mda 2
  mda-type cx2-100g-cfp
  no shutdown
exit
no shutdown
-----
```

2.18.3.1.1 Configuring MDA Access and Network Pool Parameters

MDA-level pools are used by ingress network queues. Network policies can be applied (optional) to create and edit QoS pool resources on egress network ports, channels, and ingress MDAs. Network-queue and slope policies are configured in the `config>qos` context.

The following example shows an MDA pool configuration for 7750 SR or 7450 ESS:

```
A:ALA-B>config>card>mda# info
-----
mda-type m10-lgb-xp-sfp
network
  egress
    pool
      slope-policy "B"
    exit
  exit
exit
access
  ingress
    pool
      resv-cbs 50
      slope-policy "A"
    exit
  exit
exit
exit
-----
A:ALA-B>config>card>mda#
```

The following example shows an XMA pool configuration for 7950 XRS:

```
A:ALA-B>config>card>mda# info
-----
mda-type cx20-10g-sfp
network
  egress
    pool
      slope-policy "B"
    exit
  exit
exit
access
  ingress
    pool
      resv-cbs 50
      slope-policy "A"
    exit
  exit
exit
exit
-----
A:ALA-B>config>card>mda#
```

2.18.3.1.2 Configuring MDA Policies for Named Pools Mode

Network ingress queues can use either MDA ingress named pools or ingress default pools but not port named pools. In the case with an IOM with multiple MDAs sharing the same buffer space (iom3-xp), network ingress queues will use only the MDA 1 named pools. Even if named pools are configured for MDA 2, they will not be used by network ingress queues. Network ingress queues configured to use MDA2 named pools will be considered pool orphaned. To check for orphan queues, use the command “show mda <mda> qos ingress orphaned-queues”.

SAP shared queues use by default the SAP shared pool; a system reserved buffer pool. Shared queues can be configured to use MDA named pools. Shared queues cannot be configured to use port pools since they are not port specific queues. In case a shared queue is configured to use a port named pool, the queue will be considered orphan and will get buffers from access ingress default pool.

For complete QoS configuration details, reference the Named Pools section of the Quality of Service Guide. Interface Named Pools configuration details are located in the Interface CLI portion of this guide.

2.18.3.2 Configuring Cards, MCMs, and MCAs

Card configurations must include a chassis slot designation. A slot must be preconfigured with the type of cards, MCMs, and MDAs which are allowed to be provisioned.

Output for Media Dependent Adapters (MDAs) show an “m” in the mda-type description, for example, m60-eth10/100-tx.

Use the **config > info** command to display card configuration information:

```
A:PE-1>config>card# info
-----
card-type iom4-e
mda 1
mda-type me10-10gb-sfp+
no shutdown
exit
mda 2
mda-type me1-100gb-cfp2
no shutdown
exit
no shutdown
-----
A:PE-1>config>card#
```

2.18.3.2.1 Configuring Cards and CMAs

Card configurations must include a chassis slot designation. A slot must be preconfigured with the type of cards and CMAs (Compact Media Adapters) which are allowed to be provisioned.

CMAs are configured using the MDA command. Output for Compact Media Adapter MDAs show a “c” in the mda-type description, for example, c8-10/100eth-tx.

Use the **config>info** command to display card configuration information:

```
A:PE1>config# info
. . .
#-----
echo "Card Configuration"
#-----
    card 1
        card-type iom-xp
    mda 5
        mda-type c8-10/100eth-tx
    exit
    mda 6
        mda-type c8-10/100eth-tx
    exit
    exit
#-----
```

2.18.3.3 Configuring Ports

This section provides the CLI syntax and examples to configure port parameters.

2.18.3.3.1 Configuring Port Pool Parameters

The buffer space is portioned out on a per port basis. Each port gets an amount of buffering which is its fair-share based on the port's bandwidth compared to the overall active bandwidth.

This mechanism takes the buffer space available and divides it into a portion for each port based on the port's active bandwidth relative to the amount of active bandwidth for all ports associated with the buffer space. The number of ports sharing the same buffer space depends on the type of MDAs populated on the IOM. An active port is considered to be any port that has an active queue associated. Once a queue is created for the port, the system will allocate the appropriate amount of buffer space to the port. This process is independently performed for both ingress and egress.

Normally, the amount of active bandwidth is considered as opposed to total potential bandwidth for the port when determining the port's fair share. If a port is channelized and not all bandwidth is allocated, only the bandwidth represented by the configured channels with queues configured is counted towards the bandwidth represented by the port. Also, if a port may operate at variable speeds (as in some Ethernet ports), only the current speed is considered. Based on the above, the number of buffers managed by a port may change due to queue creation and deletion, channel creation and deletion and port speed variance on the local port or other ports sharing the same buffer space.

After the active bandwidth is calculated for the port, the result may be modified through the use of the **ing-percentage-of-rate** and **egr-percent-of-rate** commands. The default value of each is 100% which allows the system to use all of the ports active bandwidth when deciding the relative amount of buffer space to allocate to the port. When the value is explicitly modified, the active bandwidth on the port is changed according to the specified percentage. If a value of 50% is given, the ports active bandwidth will be multiplied by 5, if a value of 150% is given, the active bandwidth will be multiplied by 1.5. This capability is independent of named pool mode. The ports rate percentage parameters may be modified at any time.

Examples:

1. To modify (in this example, to double) the size of buffer allocated on ingress for a port:

```
B:SR7-10# configure port 1/2/1 modify-buffer-allocation-  
rate ing-percentage-of-rate 200
```

2. To modify (in this example, to double) the size of buffer allocated on ingress for a port:

```
B:SR7-10# configure port 1/2/1 modify-buffer-allocation-  
rate egr-percentage-of-rate 200
```

The Named Buffer Pools feature provides a way to customize the port ingress and/or egress buffer allocation. The port buffer allocation size and the Forwarding class (FC) queue association to the buffer pool can be changed. By mapping each FC to different pools, it is possible to achieve separation of the available buffers per forwarding class.

Previous to this feature only the default buffer allocation mode was available, with Buffer allocation has the following characteristics:

- Each port manages a buffer according to its active bandwidth (ports with equal active bandwidth get the same buffer size).
- An access port has 2 default pools created: access-ingress and access-egress.
- A network port has 2 default pools created: ingress-MDA (common pool for all ingress network ports) and network-egress.

- All queues defined for a port receive buffers from the same buffer pool.

The Named Buffer Pools feature offers the following new capabilities:

- Ability to modify the port bandwidth considered for buffer allocation without changing the active port bandwidth. (modify-buffer-allocation-rate) (ports with equal active bandwidth can be configured to get different buffer size)
- Configuration of a named pool policy that includes the customized buffer pools
- Forwarding class queues are associated with the named pools
- Pools can be default, MDA common pools, or port specific pools

The following example shows port pool configurations:

```
A:ALA-B>config>port# info
-----
      access
        egress
          pool
            slope-policy "slopePolicy1"
          exit
        exit
      exit
    network
      egress
        pool
          slope-policy "slopePolicy2"
        exit
      exit
    exit
  no shutdown
-----
```

The following shows a CBS configuration over subscription example:

```
*A:Dut-T>config>port# info
-----
      access
        ingress
          pool
            amber-alarm-threshold 10
            resv-cbs 10 amber-alarm-action step 1 max 30
          exit
        exit
      exit
    ethernet
      mode access
      encap-type dot1q
    exit
  no shutdown
```

2.18.3.3.2 Changing Hybrid-Buffer-Allocation

The following example shows a hybrid-buffer-allocation value change (from default) for ingress. In this example, the network-egress buffer pool is two times the size of the access-egress.

```
A:SR>config>port>hybrid-buffer-allocation# info
-----
egr-weight access 20 network 40
```

2.18.3.3.3 Changing APS Parameters



Note: Nokia recommends grouping working lines and protect lines on separate IOMs.

APS configuration rules:

- A working port must be added first. Then a protection port can be added or removed at any time.
- A protection port must be shutdown before being removed from an APS group.
- A path cannot be configured on a port before the port is added to an APS group.
- A working port cannot be removed from an APS group until the APS port path is removed.
- When ports are added to an APS group, all path-level configurations are available only on the APS port level and configuration on the physical member ports are blocked.
- For a multi-chassis APS group, only one member circuit (either working or protect) can be added. Note that the neighbor IP address of an APS group must be configured before adding a member circuit in it. The configuration of a non-zero neighbor IP address indicates the APS group as multi-chassis. Thus, the member circuit and services must be removed before adding or removing the neighbor IP address (for example, before converting an APS group from multi-chassis to single-chassis or single-chassis to multi-chassis).
- Bundle Protection Group (BPGp) — A BPGp is a collection of two bundles created on the APS Group port. Working bundle resides on the working circuit of the APS group, while protection bundle resides on the protection circuit of the APS group. APS protocol running on the circuits of the APS Group port monitors the health of the Sonet/SDH line and based on it or administrative action moves user traffic from one bundle to another in the group as part of an APS switch.

The following shows a sample configuration for an ATM SC-APS group that contains an aPipe SAP:

```
A:ALA-274>config# port (1/1/1)
-----
sonet-sdh
  speed oc3
exit
no-shutdown
-----
A:ALA-274>config>port# aps-1
-----
aps
  working-circuit 1/1/1
  protect-circuit 1/1/2
exit
sonet-sdh
  path
    atm
    exit
    no-shutdown
  exit
exit
no-shutdown
exit
-----
A:ALA-274>config>service# apipe 100
-----
sap aps-1:0/100 create
exit
spoke-sdp 1:100 create
exit
no-shutdown
-----
```

The following shows an example of the configuration for the working circuit/node of an MC-APS group:

```
A:ALA-274>config>port (2/1/1)# info
-----
description "APS Group"
aps
  neighbor 192.0.2.2
  working-circuit 2/1/1
exit
no shutdown
-----
A:ALA-274>config>port#
```

The following shows an example of the configuration for the protect circuit/node of an MC-APS group:

```
A:ALA-274>config>port (2/2/2)# info
-----
description "APS Group"
aps
```

```

        neighbor 192.0.2.1
        protect-circuit 2/2/2
    exit
    no shutdown
-----
A:ALA-274>config>port#

```

2.18.3.3.4 Configuring Ethernet Port Parameters

Ethernet Network Port

A network port is network facing and participates in the service provider transport or infrastructure network processes.

The following example shows a network port configuration:

```

A:ALA-B>config>port# info
-----
        description "Ethernet network port"
        ethernet
    exit
    no shutdown
-----
A:ALA-B>config>port#

```

Ethernet Access Port

Services are configured on access ports used for customer-facing traffic. If a Service Access Port (SAP) is to be configured on a port, it must be configured as access mode. When a port is configured for access mode, the appropriate encapsulation type can be specified to distinguish the services on the port. Once a port has been configured for access mode, multiple services may be configured on the port.

The following example shows an Ethernet access port configuration:

```

A:ALA-A>config>port# info
-----
        description "Ethernet access port"
        access
            egress
                pool
                    slope-policy "slopePolicy1"
            exit
        exit
    exit
    network
        egress
            pool

```

```
        slope-policy "slopePolicy2"
      exit
    exit
  exit
  ethernet
    mode access
    encap-type dot1q
  exit
  no shutdown
-----
A:ALA-A>config>port#
```

Configuring 802.1x Authentication Port Parameters

The following example shows an 802.1x port configuration:

```
A:ALA-A>config>port>ethernet>dot1x# info detail
-----
        port-control auto
        radius-plcy dot1xpolicy
        re-authentication
        re-auth-period 3600
        max-auth-req 2
        transmit-period 30
        quiet-period 60
        supplicant-timeout 30
        server-timeout 30
        no tunneling
-----
```

2.18.3.3.5 Configuring SONET/SDH Port Parameters

SONET/SDH features can only be configured on ports on the following MDAs and CMAs:

- OC-3
- OC-3 ASAP
- OC-12/3
- OC-48
- OC-192
- OC-768
- OC-12 ASAP
- Channelized OC3
- Channelized OC12
- ATM OC-12/3

- ATM OC-12
- Channelized ASAP OC3
- Channelized ASAP OC12

When an Ethernet port is configured in WAN mode (xgig wan), you can change certain SONET/SDH parameters to reflect the SONET/SDH requirements for this port.

The following CLI output shows an example of a SONET/SDH configuration for a WAN PHY Ethernet port.

```
*A:7950 XRS-20>config>port# info
-----
shutdown
ethernet
    xgig wan
exit
sonet-sdh
    tx-dus
    suppress-lo-alarm
    threshold ber-sd rate 4
    section-trace increment-z0
    path
        trace-string "hello"
        report-alarm pais
        signal-label 0x20
    exit
exit
-----
```

SONET/SDH Network Port

The following example shows a SONET/SDH network mode configuration:

```
A:ALA-A>config>port# info
-----
description "SONET/SDH network port"
sonet-sdh
    path
        no shutdown
    exit
exit
no shutdown
-----
A:ALA-A>config>port#
```

SONET/SDH Access Port

The following example shows a SONET/SDH access port configuration for the 7750 SR:

```
A:ALA-A>config>port# info
-----
      description "SONET/SDH access port"
      sonet-sdh
        path
          mode access
          encap-type frame-relay
          mac 00:03:47:c8:b4:86
          frame-relay
          exit
          no shutdown
        exit
      exit
      no shutdown
-----
A:ALA-A>config>port#
```

2.18.3.3.6 Configuring Channelized Ports

When configuring channelized ports, the port ID is specified in different ways depending on the MDA type and level of channelization. Ethernet ports cannot be channelized. [Table 38](#) lists the channelization options and port syntax available on the 7750 SR channelized MDAs.

Table 38 Channelization Options Available on the 7750 SR Channelized MDAs

Framing	Channelization/Mapping Option	Channelized MDAs Supporting Services on the Port/ Channel
599,040 kbits/s (clear channel OC12/STM-4)		
SDH	STM4>AUG4>VC4-C4	None
SONET	OC12>STS12>STS12c SPE	None
139,264 kbits/s ñ 149,760 Kbits/s (clear channel STS-3/STM-1 or STS-3/STM-1 channel within STS12-STM4		
SDH	STM4>AUG4>AUG1>VC4	m4-choc3-as
SONET	OC12>STS12>STS3c SPE	m4-choc3-as
44,763 kbits/s (DS3 or sub-DS3 port or a channel)		

Table 38 Channelization Options Available on the 7750 SR Channelized MDAs (Continued)

Framing	Channelization/Mapping Option	Channelized MDAs Supporting Services on the Port/Channel
SDH	STM4>AUG4>AUG1>VC4>TUG3>VC3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SONET	OC12>STS12>STS1 SPE	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC4>TUG3>VC3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SONET	OC12>STS12>STS1 SPE	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
Up to 2,048 kbits/s (n*DS0 within E1 up to E1)		
SDH	STM4>AUG4>AUG1>VC4>TUG3>TUG2>VC12	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC3>TUG2>VC12	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as

Table 38 Channelization Options Available on the 7750 SR Channelized MDAs (Continued)

Framing	Channelization/Mapping Option	Channelized MDAs Supporting Services on the Port/ Channel
SDH	STM4>AUG4>AUG1>VC4>TUG3>VC3>DS3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC3>DS3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SONET	OC12>STS12>STS1 SPE>VT GROUP>VT2 SPE	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SONET	OC12>STS12>STS1 SPE>DS3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
Up to 1,544 kbits/s (n*DS0 within DS1 up to DS1)		
SDH	STM4>AUG4>AUG1>VC4>TUG3>TUG2>TU11>VC11	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC4>TUG3>TUG2>TU12>VC11	None
SDH	STM4>AUG4>AUG1>VC3>TUG2>VC11	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC4>TUG3>TUG2>VC12	m1-choc12 m4-choc3 m12-chds3

Table 38 Channelization Options Available on the 7750 SR Channelized MDAs (Continued)

Framing	Channelization/Mapping Option	Channelized MDAs Supporting Services on the Port/Channel
SDH	STM4>AUG4>AUG1>VC3>TUG2>VC12	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC4>TUG3>VC3>DS3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SDH	STM4>AUG4>AUG1>VC3>DS3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SONET	OC12>STS12>STS1 SPE>VT GROUP>VT1.5 SPE	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as
SONET	OC12>STS12>STS1 SPE>VT GROUP>VT2 SPE	m1-choc12 m4-choc3 m12-chds3
SONET	OC12>STS12>STS1 SPE>DS3	m1-choc12 m4-choc3 m12-chds3 m4-choc3-as



Note: The E1 encapsulation in the ASAP MDA and in the channelized MDAs is compliant to G.704 and G.703. The G.703 feature allows a user to configure an unstructured E1 channel on deep channel MDAs and ASAP MDAs. In G.704, time slot 0 carries timing information by a service provider and thus, only 31 slots are available to the end user. In G.703, all 32 time slots are available to the end user. Timing is provided by the end user.

A port ID for channels has one of the following syntax as applicable to channelization and mapping options where the port configuration syntax is slot/mda/port ([Table 39](#)).

Table 39 Channelized Port Syntax Examples

Port ID for Physical Port Speed			
Channel speed	OC12/STM4	OC3/STM1	DS3/E3
SONET/SDH			
STS12/STM4	port.sts12	N/A	N/A
STS3/STM1	port.sts3-{1 to 4}	port.sts3	N/A
STS1/STM0	port.sts1-{1 to 4}. {1 to 3}	port.sts1-{1 to 3}	N/A
TUG3	port.tug3-{1 to 4}. {1 to 3}	port.tug3-{1 to 3}	N/A
TU3	port.tu3-{1 to 4}. {1 to 3}	port.tu3-{1 to 3}	N/A
VT15/VC1.1 ¹	port.vt15-{1 to 4}. {1 to 3}. {1 to 4}. {1 to 7}	port.vt15-{1 to 3}. {1 to 4}. {1 to 7}	N/A
VT2/VC12 ¹	port.vt2-{1 to 4}. {1 to 3}. {1 to 3}. {1 to 7}	port.vt2-{1 to 3}. {1 to 3}. {1 to 7}	N/A
TDM			
DS3/E3	port.{1 to 4}. {1 to 3}	port.{1 to 3}	port
DS1 in DS3	port.{1 to 4}. {1 to 3}. {1 to 28}	port.{1 to 3}. {1 to 28}	port.{1 to 28}
DS1 in VT2	port.{1 to 4}. {1 to 3}. {1 to 3}. {1 to 7}	port.{1 to 3}. {1 to 3}. {1 to 7}	N/A
DS1 in VT15	port.{1 to 4}. {1 to 3}. {1 to 4}. {1 to 7}	port.{1 to 3}. {1 to 4}. {1 to 7}	N/A
E1 in DS3	port.{1 to 4}. {1 to 3}. {1 to 21}	port.{1 to 3}. {1 to 21}	port.{1 to 21}
E1 in VT2	port.{1 to 4}. {1 to 3}. {1 to 3}. {1 to 7}	port.{1 to 3}. {1 to 3}. {1 to 7}	N/A
N*DS0 in DS1 in DS3	port.{1 to 4}. {1 to 3}. {1 to 28}. {1 to 24}	port.{1 to 3}. {1 to 28}. {1 to 24}	port.{1 to 28}. {1 to 24}
N*DS0 in DS1 in VT2	port.{1 to 4}. {1 to 3}. {1 to 3}. {1 to 7}. {1 to 24}	port.{1 to 3}. {1 to 3}. {1 to 7}. {1 to 24}	N/A
N*DS0 in DS1 in VT15	port.{1 to 4}. {1 to 3}. {1 to 4}. {1 to 7}. {1 to 24}	port.{1 to 3}. {1 to 4}. {1 to 7}. {1 to 24}	N/A
N*DS0 in E1 in DS3	port.{1 to 4}. {1 to 3}. {1 to 21}. {2 to 32}	port.{1 to 3}. {1 to 21}. {2 to 32}	port.{1 to 21}. {2 to 32}
N*DS0 in E1 in VT2	port.{1 to 4}. {1 to 3}. {1 to 3}. {1 to 7}. {2 to 32}	port.{1 to 3}. {1 to 3}. {1 to 7}. {2 to 32}	N/A

Note:

1. Supported by TDM satellite.

Verify the MDA Type


To ensure that you have a channel-capable MDA, verify that the MDA-type you are configuring by entering a **show mda slot-id** command.

The MDAs shown in the *MDA Provisioned* column in the following output are a 12-port channelized DS3 MDA (m12-ds3) on card 1, MDA slot 1, and a 1-port channelized OC12-SFP MDA (m1-choc12-sfp) on card 1, MDA slot 2.

```
A:ALA-A# show mda 1
=====
MDA 1/1
=====
Slot  Mda  Provisioned      Equipped      Admin      Operational
      Mda      Mda-type          Mda-type          State      State
-----
1      1      m12-ds3           m12-ds3           up         provisioned
=====
ALA-A# show mda 2
=====
MDA 1/2
=====
Slot  Mda  Provisioned      Equipped      Admin      Operational
      Mda      Mda-type          Mda-type          State      State
-----
1      2      m1-choc12-sfp     m1-choc12-sfp     up         provisioned
=====
A:ALA-A#
```

Configuring a Channelized DS3 Port

[Figure 49](#) shows the logic of the DS3 port configuration.

Figure 49 Channelized DS3 Port Structure

The following shows the steps to configure a channelized port on a 12-port DS3 MDA:

```
A:ALA-A>config# port 7/1/1
A:ALA-A>config>port# tdm
```

In order to set the channelized mode on a port, the DS3 parameter must be in a shut down state. Clear channel uses out-of-band signaling, not in-band signaling, so the channel's entire bit rate is available. Channelized ports use in-band signaling and must be explicitly enabled as shown:

```
A:ALA-A>config>port>tdm# ds3
A:ALA-A>config>port>tdm>ds3# shutdown
A:ALA-A>config>port>tdm>ds3# channelized ds1
A:ALA-A>config>port>tdm>ds3# no shutdown
A:ALA-A>config>port>tdm>ds3# exit
```

In the DS1 context, configure DS0 channel groups parameters. 24 timeslots can be configured per channel group as shown:

```
A:ALA-A>config>port>tdm# ds1 1
A:ALA-A>config>port>tdm>ds1# no shutdown
A:ALA-A>config>port>tdm>ds1# channel-group 1
A:ALA-A>config>port>tdm>ds1>channel-group# timeslots 1
A:ALA-A>config>port>tdm>ds1>channel-group# encap-type frame-relay
A:ALA-A>config>port>tdm>ds1>channel-group# no shutdown
A:ALA-A>config>port>tdm>ds1>channel-group# exit
A:ALA-A>config>port>tdm>ds1# channel-group 2
A:ALA-A>config>port>tdm>ds1>channel-group# timeslots 2-10
A:ALA-A>config>port>tdm>ds1>channel-group# no shutdown
```

```
A:ALA-A>config>port>tdm>ds1>channel-group# exit
A:ALA-A>config>port>tdm>ds1# exit
A:ALA-A>config>port>tdm# ds1 2
A:ALA-A>config>port>tdm>ds1# channel-group 1
A:ALA-A>config>port>tdm>ds1>channel-group# timeslots 1
A:ALA-A>config>port>tdm>ds1>channel-group# exit
A:ALA-A>config>port>tdm>ds1# no shutdown
A:ALA-A>config>port>tdm>ds1# channel-group 2
A:ALA-A>config>port>tdm>ds1>channel-group# timeslots 2
A:ALA-A>config>port>tdm>ds1>channel-group# exit
A:ALA-A>config>port>tdm>ds1# no shutdown
```

The following output shows the channelized mode configuration:

```
A:ALA-A>config>port># info
-----
      tdm
        ds3 ds3
          channelized ds1
          no shutdown
        exit
      ds1 ds1-1
        channel-group 1
          encap-type frame-relay
          timeslots 1
          frame-relay
          exit
          no shutdown
        exit
        channel-group 2
          shutdown
          timeslots 2-10
        exit
        no shutdown
      exit
      ds1 ds1-2
        channel-group 1
          shutdown
          timeslots 1
        exit
        channel-group 2
          timeslots 2
          no shutdown
        exit
        no shutdown
      exit
    exit
    no shutdown
-----
A:ALA-A>config>port#
```

Services can be applied to the configured channelized ports. The following example shows the CLI usage to configure a customer IES service with interface SAPs on the channelized ports. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for information about how to configure services.

```
A:ALA-A>config>service# ies 103 customer 1 create
A:ALA-A>config>service>ies$ interface test1 create
A:ALA-A>config>service>ies>if$ address 192.168.1.1/24
A:ALA-A>config>service>ies>if# sap 7/1/1.1.2 create
A:ALA-A>config>service>ies>if>sap$ exit
A:ALA-A>config>service>ies>if# no shutdown
A:ALA-A>config>service>ies>if# exit
A:ALA-A>config>service>ies# interface test2 create
A:ALA-A>config>service>ies>if$ address 192.168.2.1/24
A:ALA-A>config>service>ies>if$ sap 7/1/1.2.1 create
A:ALA-A>config>service>ies>if>sap$ exit
A:ALA-A>config>service>ies>if# no shutdown
A:ALA-A>config>service>ies>if# exit
A:ALA-A>config>service>ies>if#
```


The following output shows the channelized ports (7/1/1.1.1 and 7/1/1.1.2) applied to SAPs on the IES service configuration:

```
A:ALA-A>config>service>ies# info
-----
...
    ies 103 customer 1 vpn 103 create
        interface "test2" create
            address 192.168.2.1/24
            sap 7/1/1.2.1 create
            exit
        exit
        interface "test1" create
            address 192.168.1.1/24
            sap 7/1/1.1.2 create
            exit
        exit
        no shutdown
    exit
...
-----
A:ALA-A>config>service>ies#
```

Configuring a Channelized OC-12-SFP Port

[Figure 50](#) shows the logic of the channelized OC-12 port configuration.

Figure 50 Channelized OC-12 Port Structure



The following shows an example to configure a channelized port on a 1-port channelized OC-12-SFP MDA:

```
ALA-A>config# port 5/2/1
```

At this level you must choose the tributary. When provisioning DS3 nodes on a channelized OC-12 MDA, you must provision the parent STS1-1 SONET path first.

```
A:ALA-A>config>port# sonet-sdh
A:ALA-A>config>port>sonet-sdh# path sts1-1.1
A:ALA-A>config>port>sonet-sdh>path# no shutdown
A:ALA-A>config>port>sonet-sdh>path# exit
```

The following shows the output:

```
A:ALA-A>config>port>sonet-sdh# info
-----
sonet-sdh
  path sts1-1.1
    no shutdown
  exit
exit
-----
A:ALA-A>config>port>sonet-sdh#
```

In order to set the channelized mode on a port, the DS3 parameter must be in a shut down state. Clear channel uses out-of-band signaling, not in-band signaling, so the channel's entire bit rate is available. Channelized ports use in-band signaling and must be explicitly enabled.

```
A:ALA-A>config>port# tdm
A:ALA-A>config>port>tdm# ds3 1.1
A:ALA-A>config>port>tdm>ds3# shutdown
A:ALA-A>config>port>tdm>ds3# channelized ds1
A:ALA-A>config>port>tdm>ds3# no shutdown
A:ALA-A>config>port>tdm>ds3# exit
```

The following shows an example of the output:

```
A:ALA-A>config>port# info
-----
sonet-sdh
  path sts12
    no shutdown
  exit
  path sts3-1
    no shutdown
  exit
  path sts1-1.1
    no shutdown
  exit
exit
tdm
  ds3 ds3-1.1
    channelized
    no shutdown
  exit
exit
no shutdown
-----
A:ALA-A>config>port#
```

In the TDM context, configure DS0 channel groups parameters. 24 timeslots can be configured per channel group.

```
A:ALA-A>config>port>tdm# ds1 1.1.1
A:ALA-A>config>port>tdm>ds1# no shutdown
A:ALA-A>config>port>tdm>ds1# channel-group 1
A:ALA-A>config>port>tdm>ds1>channel-group# timeslots 1
A:ALA-A>config>port>tdm>ds1>channel-group# no shutdown
A:ALA-A>config>port>tdm>ds1>channel-group# exit
A:ALA-A>config>port>tdm>ds1# no shutdown
A:ALA-A>config>port>tdm>ds1# channel-group 2
A:ALA-A>config>port>tdm>tds1>channel-group# timeslots 2
A:ALA-A>config>port>tdm>ds1>channel-group# no shutdown
A:ALA-A>config>port>tdm>ds1>channel-group# exit
A:ALA-A>config>port>tdm>ds1# exit
```

```
A:ALA-A>config>port>tdm# info
-----
```

```

sonet-sdh
  path sts12
    no shutdown
  exit
  path sts3-1
    no shutdown
  exit
  path sts1-1.1
    no shutdown
  exit
exit
tdm
  ds3 ds3-1.1
    channelized
    no shutdown
  exit
  ds1 ds1-1.1.1
    channel-group 1          (see SAP 5/2/1.1.1.1.1 below)
    timeslots 1
    no shutdown
  exit
  channel-group 2          (see SAP 5/2/1.1.1.1.2 below)
    timeslots 2
    no shutdown
  exit
  no shutdown
exit
no shutdown
-----
A:ALA-A>config>port>tdm#

```

Services can be applied to the configured channelized ports. The following example shows the CLI usage to configure a customer IES service with interface SAPs on the channelized ports. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for information about how to configure services.

```

A:ALA-A>config>service# ies 104 customer 1 create
A:ALA-A>config>service>ies$ interface testA create
A:ALA-A>config>service>ies>if$ address 192.168.1.1/24
A:ALA-A>config>service>ies>if# sap 5/2/1.1.1.1.1 create
A:ALA-A>config>service>ies>if>sap$ exit
A:ALA-A>config>service>ies>if# no shutdown
A:ALA-A>config>service>ies>if# exit
A:ALA-A>config>service>ies# interface testB create
A:ALA-A>config>service>ies>if$ address 192.168.2.1/24
A:ALA-A>config>service>ies>if# sap 5/2/1.1.1.1.2 create
A:ALA-A>config>service>ies>if>sap$ exit
A:ALA-A>config>service>ies>if# no shutdown
A:ALA-A>config>service>ies>if# exit
A:ALA-A>config>service>ies# no shutdown

```

The following output shows the channelized ports 5/2/1.1.1.1.1 and 5/2/1.1.1.1.2) applied to SAPs on the IES service configuration:

```

A:ALA-A>config>service>ies# info

```

```

-----
        interface "testA" create
            address 192.168.1.1/24
            sap 5/2/1.1.1.1.1 create
            exit
        exit
        interface "testB" create
            address 192.168.2.1/24
            sap 5/2/1.1.1.1.2 create
            exit
        exit
        no shutdown
-----
A:ALA-A>config>service>ies#

```

Configuring a Channelized Any Service Any Port (ASAP) OC3-SFP Port

This section provides examples to configure PPP, FR, cHDLC, and ATM n*DS0 channels on a channelized port on channelized ASAP OC-3 SFP MDA in slot 1/1/1. The ASAP OC-12 SFP MDA also supports the SONET options.

```
ALA-A>config# port 1/1/1
```

At this level you must choose the tributary. When provisioning DS3 nodes on a channelized ASAP OC-3 MDA, you must provision the parent STS1-1 SONET path first.

```

A:ALA-A>config>port# sonet-sdh
A:ALA-A>config>port>sonet-sdh# framing sdh
A:ALA-A>config>port>sonet-sdh# path sts1-1
A:ALA-A>config>port>sonet-sdh>path# no shutdown
A:ALA-A>config>port>sonet-sdh>path# exit
A:ALA-A>config>port>sonet-sdh# info
-----
        sonet-sdh
            framing sdh
            path sts1-1
                no shutdown
            exit
        exit
-----
A:ALA-A>config>port>sonet-sdh#

```

In order to set the channelized mode on a port, the DS3 parameter must be in a shut down state. Clear channel uses out-of-band signaling, not in-band signaling, so the channel's entire bit rate is available. Channelized ports use in-band signaling and must be explicitly enabled.

```

A:ALA-A>config>port# tdm
A:ALA-A>config>port>tdm# ds3 1
A:ALA-A>config>port>tdm>ds3# shutdown
A:ALA-A>config>port>tdm>ds3# channelized e1

```

```
A:ALA-A>config>port>tdm>ds3# no shutdown
A:ALA-A>config>port>tdm>ds3# exit
A:ALA-A>config>port# info
-----
sonet-sdh
    path sts1-1
        no shutdown
    exit
exit
tdm
    ds3 1
        channelized e1
        no shutdown
    exit
exit
no shutdown
-----
A:ALA-A>config>port#
```

In the TDM E1 context, configure DS0 channel groups and their parameters. For a DS1 channel-group, up to 24 timeslots can be assigned (numbered 1 to 24). For an E1 channel-group, up to 31 timeslots can be assigned (numbered 2 to 32). For ATM, all timeslots are auto-configured when a channel group is created (there is no sub-E1 for ATM). ATM, Frame Relay and BCP-NULl encapsulation examples follow:

```
A:ALA-A>config>port>tdm# e1 1.1
A:ALA-A>config>port>tdm>e1# channel-group 1
A:ALA-A>config>port>tdm>e1>channel-group# timeslots 2
A:ALA-A>config>port>tdm>e1>channel-group# no shutdown
A:ALA-A>config>port>tdm>e1>channel-group#
A:ALA-A>config>port>tdm>e1# no shutdown
A:ALA-A>config>port>tdm>e1# channel-group 2
A:ALA-A>config>port>tdm>e1>channel-group# timeslots 3
A:ALA-A>config>port>tdm>e1>channel-group# encap-type frame-relay
A:ALA-A>config>port>tdm>e1>channel-group# no shutdown
A:ALA-A>config>port>tdm>e1>channel-group# exit
A:ALA-A>config>port>tdm>e1# channel-group 3
A:ALA-A>config>port>tdm>e1>channel-group# timeslots 11,12
A:ALA-A>config>port>tdm>e1>channel-group# encap-type cisco-hdlc
A:ALA-A>config>port>tdm>e1>channel-group# no shutdown
A:ALA-A>config>port>tdm>e1>channel-group# exit
A:ALA-A>config>port>tdm>e1# no shutdown
A:ALA-A>config>port>tdm>e1# exit
A:ALA-A>config>port>tdm# e1 1.2
A:ALA-A>config>port>tdm>e1# no shutdown
A:ALA-A>config>port>tdm>e1# channel-group 1
A:ALA-A>config>port>tdm>e1>channel-group# encap-type atm
A:ALA-A>config>port>tdm>e1>channel-group# no shutdown
A:ALA-A>config>port>tdm>e1>channel-group# exit
A:ALA-A>config>port>tdm>e1# no shutdown
A:ALA-A>config>port>tdm# info
-----
tdm
    ds3 1
        channelized e1
        no shutdown
    exit
```

```

e1 1.1
  channel-group 1
    timeslots 2
    no shutdown
  exit
  channel-group 2
    encap-type frame-relay
    frame-relay
    exit
    timeslots 10
    no shutdown
  exit
  channel-group 3
    encap-type cisco-hdlc
    cisco-hdlc
    exit
    timeslots 11,12
    no shutdown
  exit
  no shutdown
exit
e1 1.2
  channel-group 1
    encap-type atm
    atm
    exit
    no shutdown
  exit
  no shutdown
exit
no shutdown
-----
A:ALA-A>config>port>tdm#

```

Services can now be applied to the configured channelized ports. Follow examples of other channelized ports in this document.

Configuring Cisco HDLC on a Channelized Port

Use the following CLI syntax to configure cHDLC:

CLI Syntax:

```

config# port port-id
tdm
    ds3 [sonet-sdh-index]
        channelized {ds1|e1}
        no shutdown
    ds1
        channel-group channel-group
        cisco-hdlc
        down-count down-count
        keepalive time-interval
        up-count up-count

```

```

encap-type {bcp-null|bcp-dot1q|ipcp|ppp-
            auto|frame-relay|wan-mirror|cisco-
            hdlc}
timeslots timeslots
no shutdown

```

The following example shows SONET/SDH access mode configuration command usage:

Example:

```

A:ALA-29>config>port>tdm# ds3
A:ALA-29>config>port>tdm>ds3# channelized ds1
A:ALA-29>config>port>tdm>ds3# no shutdown
A:ALA-29>config>port>tdm>ds3# exit
A:ALA-29>config>port>tdm# ds1 1
A:ALA-29>config>port>tdm>ds1# no shutdown
A:ALA-29>config>port>tdm>ds1# channel-group 1
A:ALA-29>config>port>tdm>ds1>channel-group# timeslots 1-
20
A:ALA-29>config>port>tdm>ds1>channel-group# encap-type
cisco-hdlc
A:ALA-29>config>port>tdm>ds1>channel-group# exit
A:ALA-29>config>port>tdm>ds1# channel-group 1
A:ALA-29>config>port>tdm>ds1>channel-group# no shutdown
A:ALA-29>config>port>tdm>ds1>channel-group# exit
A:ALA-29>config>port>tdm>ds1# exit
A:ALA-29>config>port>tdm#

```

The following example shows a configuration:

```

A:ALA-29>config>port# inf
-----
      tdm
        ds3
          channelized ds1
          no shutdown
        exit
      ds1 1
        channel-group 1
          encap-type cisco-hdlc
          timeslots 1-20
          cisco-hdlc
          exit
          no shutdown
        exit
      exit
    exit
  no shutdown
-----
A:ALA-29>config>port#

```

2.18.3.3.7 Configuring Channelized STM1/OC3 Parameters

The following example shows basic syntax to configure channelized STM1/OC3 parameters:

CLI Syntax:

```
config# port port-id
sonet-sdh
    framing {sonet | sdh}
    group sonet-sdh-index payload {tu3 | vt2 | vt15}
    path [sonet-sdh-index]
        payload {sts3 | tug3 | ds3 | e3}
        trace-string [trace-string]
    no shutdown
```

Example:

```
config# port 5/2/1
config>port# sonet-sdh
config>port>sonet-sdh# framing sdh
config>port>sonet-sdh# path sts3
config>port>sonet-sdh>path# trace-string "HO-path"
config>port>sonet-sdh>path# exit
config>port>sonet-sdh# group tug3-1 payload vt2
config>port>sonet-sdh# group tug3-3 payload vt2
config>port>sonet-sdh# path vt2-1.1.1
config>port>sonet-sdh>path# trace-string "LO-path 3.7.3"
config>port>sonet-sdh>path# no shutdown
config>port>sonet-sdh>path# exit
config>port>sonet-sdh# exit
config>port# tdm
config>port>tdm# e1 1.1.1
config>port>tdm>e1# channel-group 1
config>port>tdm>e1>channel-group# timeslots 2-32
config>port>tdm>e1>channel-group# no shutdown
config>port>tdm>e1>channel-group# exit
config>port>tdm# e1 3.7.3
config>port>tdm>e1# channel-group 2
config>port>tdm>e1>channel-group# timeslots 2-32
config>port>tdm>e1>channel-group# no shutdown
config>port>tdm>e1>channel-group# exit
```

The following shows the configuration output:

```
A:ALA-49>config>port# info
-----
sonet-sdh
  framing sdh
  path sts3
    trace-string "HO-path"
  no shutdown
exit
group tug3-1 payload vt2
```



```

        group tug3-3 payload vt2
        path vt2-1.1.1
            trace-string "LO-path 3.7.3"
            no shutdown
        exit
        path vt2-3.7.3
            no shutdown
        exit
    exit
    tdm
        e1 1.1.1
            channel-group 1
            timeslots 2-32
            no shutdown
        exit
        no shutdown
    exit
    e1 3.7.3
        channel-group 2
        timeslots 2-32
        no shutdown
    exit
    no shutdown
    exit
    no shutdown
-----
A:ALA-49>config>port#

```

Sample Cpipe Port Configurations

Before a Cpipe service can be provisioned, the following entities must be configured:

Configuring a DS1 Port

The following shows an example of a DS1 port configured for CES.

```

A:sim216# show port 1/5/1.1.3.1
=====
TDM DS1 Interface
=====
Description          : DS1
Interface            : 1/5/1.1.3,1
Type                 : ds1
Admin Status         : up
Physical Link        : yes
Signal Mode          : none
Last State Change    : 10/31/2006 14:23:12
Loopback             : none
Remote Loop respond  : false
Load-balance-algo    : default
BERT Duration        : N/A
BERT Synched         : 00h00m00s
Framing              : esf
Oper Status          : up
Clock Source         : loop-timed
Channel IfIndex      : 580943939
Invert Data          : false
In Remote Loop       : false
Egr. Sched. Pol      : n/a
BERT Pattern         : none
Err Insertion Rate   : 0

```

```

BERT Errors      : 0                      BERT Status      : idle
BERT Total Bits  : 0
Cfg Alarm        : ais los
Alarm Status     :
=====
A:sim216#

```

Configuring a Channel Group

The following shows an example of a DS1 channel group configured for CES.

```

A:sim216# show port 1/5/1.1.3.1
=====
TDM DS0 Chan Group
=====
Description      : DS0GRP
Interface        : 1/5/1.1.3.1
TimeSlots        : 1-12
Speed            : 64                      CRC              : 16
Admin Status     : up                     Oper Status        : up
Last State Change : 10/31/2006 14:23:12   Chan-Grp IfIndex   : 580943940
Configured mode   : access                 Encap Type         : cem
Admin MTU         : 4112                   Oper MTU           : 4112
Physical Link     : Yes                     Bundle Number      : none
Idle Cycle Flags  : flags                  Load-balance-algo  : default
Egr. Sched. Pol   : n/a
=====
A:sim216#

```

2.18.3.3.8 Configuring ATM SAPs

ATM SAP in an IES Service

The following shows a sample IES service SAP configuration:

```

:ALA-701>config>service>ies# info
-----
      interface "atm_1" create
      address 192.168.4.1/24
      sap 2/1/1:17/24 create
      exit
    exit
      interface "atm_2" create
      address 192.168.5.1/24
      sap 2/1/1:18/300 create
      exit
    exit
      no shutdown
-----
B:ALA-701>config>service>ies#

```

ATM SAP in an Epipe Service

The following shows a sample Epipe service SAP configuration:

```
B:ALA-701>config>service# info
-----
...
        epipe 5 customer 1 create
            shutdown
            sap 2/1/2:15/25 create
            exit
            sap 2/1/3:25/35 create
            exit
        exit
-----
B:ALA-701>config>service#
```

2.18.3.3.9 Configuring DWDM Port Parameters

The following shows a sample DWDM port configuration:

```
*A:ALA-A>config>port>dwdm># info
-----
        channel 44
        wavetracker
            power-control
                target-power -7.50
            exit
            encode key1 205 key2 749
        exit
-----

*A:ALA-A>config>port>dwdm># info detail
-----
        channel 44
        wavetracker
            power-control
                target-power -7.50
            exit
            encode key1 205 key2 749
            report-alarm enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
        exit
        rxdtv-adjust
-----

*A:ALA-A>config>port>dwdm># wavetracker

*A:ALA-A>config>port>dwdm>wavetracker># info
-----
        power-control
            target-power -7.50
        exit
        encode key1 205 key2 749
-----
```

```
*A:ALA-A>config>port>dwdm>wavetracker># info detail
-----
power-control
    target-power -7.50
exit
encode key1 205 key2 749
report-alarm enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
-----
```

2.18.3.3.10 Configuring WaveTracker Parameters

The following example shows the default configuration with WaveTracker disabled:

```
*A:ALA-A>config>port>dwdm># info
-----
channel 44
-----

*A:ALA-A>config>port>dwdm># info detail
-----
channel 44
wavetracker
    no power-control
    no encode
    report-alarm enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
exit
rxdtv-adjust
-----
```

The following example shows a configuration with DWDM channel 44, WaveTracker power control transmit power at -7.5 dBm and WaveTracker encoded keys 205 and 749.

```
*A:ALA-A>config>port>dwdm># info
-----
channel 44
wavetracker
    power-control
        target-power -7.50
    exit
    encode key1 205 key2 749
exit
-----

*A:ALA-A>config>port>dwdm># info detail
-----
channel 44
wavetracker
    power-control
        target-power -7.50
    exit
    encode key1 205 key2 749
    report-alarm enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
-----
```

```

exit
rxdtv-adjust
-----
*A:ALA-A>config>port>dwdm># wavetracker

*A:ALA-A>config>port>dwdm>wavetracker># info
-----
power-control
    target-power -7.50
exit
encode key1 205 key2 749
-----

*A:ALA-A>config>port>dwdm>wavetracker># info detail
-----
power-control
    target-power -7.50
exit
encode key1 205 key2 749
report-alarm enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
-----

```

The following is an example of the show port <portId> wavetracker command for the non-default WaveTracker configuration as shown above:

```

*A:ALA-A# show port 3/2/1 wavetracker

=====
Wavelength Tracker
=====
Power Control      : Enabled                WaveKey Status    : Enabled
Target Power      : -7.50 dBm                WaveKey 1         : 205
Measured Power     : -7.49 dBm                WaveKey 2         : 749

Cfg Alarms         : enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
Alarm Status       :

Maximum Power      : 0.47 dBm                Power Upper Margin : 7.96 dB
Minimum Power      : -21.23 dBm               Power Lower Margin : 13.74 dB
=====

```

The following example shows the Wavetracker keys allowed for each DWDM channel:

ITU Channel	Key1 Min	Key1 Max	Key2 Min	Key2 Max
61	1548	1548	2032	2032
59	1	15	545	559
58	18	32	562	576
57	35	49	579	593
56	52	66	596	610
54	69	83	613	627
53	86	100	630	644
52	103	117	647	661
51	120	134	664	678

49	137	151	681	698
48	154	168	698	712
47	171	185	715	729
46	188	202	732	746
44	205	219	749	763
43	222	236	766	780
42	239	253	783	797
41	256	270	800	814
39	273	287	817	831
38	290	304	834	848
37	307	321	851	865
36	324	338	868	882
34	341	355	885	899
33	358	372	902	916
32	375	389	919	933
31	392	406	936	950
29	409	423	953	967
28	426	440	970	984
27	443	457	987	1001
26	460	474	1004	1018
24	477	491	1021	1035
23	494	508	1038	1052
22	511	525	1055	1069
21	528	542	1072	1086
60	1089	1103	1573	1587
55	1106	1120	1590	1604
50	1123	1137	1607	1621
45	1140	1154	1624	1638
40	1157	1171	1641	1655
35	1174	1188	1658	1672
30	1191	1205	1675	1689
25	1208	1222	1692	1706
20	1225	1239	1709	1723
19	1242	1256	1726	1740
18	1259	1273	1743	1757
17	1276	1290	1760	1774
595	1293	1307	1777	1791
585	1310	1324	1794	1808
575	1327	1341	1811	1825
565	1344	1358	1828	1842
545	1361	1375	1845	1859
535	1378	1392	1862	1876
525	1395	1409	1879	1893
515	1412	1426	1896	1910
495	1429	1443	1913	1927
485	1446	1460	1930	1944
475	1463	1477	1947	1961
465	1480	1494	1964	1978
445	1497	1511	1981	1995
435	1514	1528	1998	2012
425	1531	1545	2015	2029
415	1548	1562	2032	2046
395	3585	3599	2049	2063
385	3602	3616	2066	2080
375	3619	3633	2083	2097
365	3636	3650	2100	2114
345	3653	3667	2117	2131
335	3670	3684	2134	2148
325	3687	3701	2151	2165

315	3704	3718	2168	2182
295	3721	3735	2185	2199
285	3738	3752	2202	2216
275	3755	3769	2219	2233
265	3772	3786	2236	2250
245	3789	3803	2253	2267
235	3806	3820	2270	2284
225	3823	3837	2287	2301
215	3840	3854	2304	2318
605	3857	3871	2321	2335
555	3874	3888	2338	2352
505	3891	3905	2355	2369
455	3908	3922	2372	2386
405	3434	3448	3946	3960
355	3451	3465	3963	3977
305	3468	3482	3980	3994
255	3485	3499	3997	4011
205	3502	3516	4014	4028
195	3519	3533	4031	4045
185	3536	3550	4048	4062
175	3553	3567	4065	4079

2.18.3.3.11 Configuring OTU Port Parameters

The following example shows an OTU port configuration:

```
*A:ALA-A>config>port>otu# info detail
-----
    otu2-lan-data-rate 11.049
    sf-sd-method fec
    sf-threshold 5
    sd-threshold 7
    fec enhanced
    no report-alarm otu-ais otu-ber-sd otu-tim otu-iae otu-biae fec-sd
    no report-alarm fec-fail fec-uncorr odu-ais odu-oci odu-lck odu-bdi
    no report-alarm odu-tim opu-tim opu-plm
    report-alarm loc los lof lom otu-ber-sf otu-bdi fec-sf
    sm-tti
        tx auto-generated
        expected auto-generated
        no mismatch-reaction
    exit
    pm-tti
        tx auto-generated
        expected auto-generated
        no mismatch-reaction
    exit
    psi-tti
        tx auto-generated
        expected auto-generated
        no mismatch-reaction
    exit
    psi-payload
        tx auto
        expected auto
        no mismatch-reaction
```

```
exit
```

The following example shows the show port <portId> otu detail for the default OTU configuration as shown above:

```
*A:ALA-A# show port 3/2/1 otu detail
```

```
=====
OTU Interface
=====
OTU Status      : Enabled          FEC Mode      : enhanced
Async Mapping   : Disabled        Data Rate     : 11.049 Gb/s

Cfg Alarms      : loc los lof lom otu-ber-sf otu-bdi fec-sf
Alarm Status    :
SF/SD Method    : FEC              SF Threshold   : 1E-5
                                      SD Threshold    : 1E-7

SM-TTI Tx (auto) : ALA-A:3/2/1/C44
SM-TTI Ex (bytes) : (Not Specified)
SM-TTI Rx        : ALA-A:5/2/1/C34
OTU-TIM reaction  : none

PM-TTI Tx (auto) : ALA-A:3/2/1/C44
PM-TTI Ex (bytes) : (Not Specified)
PM-TTI Rx        : ALA-A:5/2/1/C34
ODU-TIM reaction  : none

PSI-TTI Tx (auto) : ALA-A:3/2/1/C44
PSI-TTI Ex (bytes) : (Not Specified)
PSI-TTI Rx        : ALA-A:5/2/1/C34
OPU-TIM reaction  : none

PSI-PT Tx (auto) : 0x03 (syncCbr)
PSI-PT Ex (auto) : 0x03 (syncCbr)
PSI-PT Rx        : 0x03 (syncCbr)
OPU-PLM reaction  : none
=====
OTU Statistics
=====
Elapsed Seconds                                10
-----
Near End Statistics                            Count
-----
FEC Corrected 0s                               0
FEC Corrected 1s                               0
FEC Unrectable Sub-rows                        0
FEC ES                                          0
FEC SES                                         0
FEC UAS                                         0
Pre-FEC BER                                   0.000E+00
Post-FEC BER                                  0.000E+00
-----
SM BIP8                                         0
SM ES                                           0
SM SES                                          0
SM UAS                                          0
```



```

SM-BIP8-BER                                0.000E+00
-----
PM BIP8                                    0
PM ES                                      0
PM SES                                    0
PM UAS                                    0
PM-BIP8-BER                                0.000E+00
-----
NPJ                                        0
PPJ                                        0
-----
Far End Statistics                          Count
-----
SM BEI                                    0
PM BEI                                    0
=====

```

The window over which the Bit Error Rate (BER) determined is based on the configured threshold level. The higher the error rate the shorter the window and as the error rate decreases the window increases. [Table 40](#) lists the configured BER thresholds and corresponding window lengths.

Table 40 Configured BER Thresholds and Window Lengths

Configured BER Threshold	Window Length
10 ⁻³	8ms
10 ⁻⁴	8ms
10 ⁻⁵	8ms
10 ⁻⁶	13ms
10 ⁻⁷	100ms
10 ⁻⁸	333ms
10 ⁻⁹	1.66s

2.18.3.3.12 Configuring ATM Interface Parameters

ATM interface parameters can only be configured for SONET/SDH ports/paths and TDM ports/channels supporting ATM encapsulation, and for IMA multilink bundles.

ATM interface parameters allow users to configure characteristics of an ATM interface. The Nokia routers support configuration of the following ATM interface characteristics:

- Cell-format — Allows user to select the ATM cell format to be used on a given interface: UNI/NNI
- ILMI — Allows user to enable/disable ILMI protocol
- Traffic-desc — Allows user to configure ILMI PVCC TM characteristics over a given ATM interface ingress and egress direction characteristics can be configured separately)
- Mapping — Allows user to select ATM cell mapping into an HDLC frame: Direct/PLCP

PLCP/Direct Mapping

Setting mapping to PLCP changes the effective speed of a DS3 interface to 40.704 M. When a port operates in a PLCP mode, the OCD events and LCD are not applicable (including related status fields and counters).

Similarly the below-defined PLCP statuses, alarms, and counters do not apply for direct mapped ports.

When a path operates in the PLCP mode, the router supports the standard ATM MIB monitoring of the PLCP operations, for example:

- PLCP severely errored framing seconds
- PLCP alarm state
- PLCP unavailable seconds counter

[Table 41](#) shows how SONET alarm status, path operational status, ATM interface and PLCP status and PLCP Alarm state interact.

Table 41 Alarm State Interactions

Content of the Received Signal						Status Field Values			
Local Signal	Local Frame	Local Payld	Local PLCP Framing	Far End Framing	Far End PLCP Framing	Path Sonet Alarm Status	Path Oper Status	Atm Interface Oper Status	PLCP Alarm State
Y	Y	Y	Y	Y	Y	None	Up	Up	No Alarm
Y	Y	Y	Y	Y	Prob	None	Up	Lower Layer Down	Far End Alarm Rx

Table 41 Alarm State Interactions (Continued)

Content of the Received Signal						Status Field Values			
Y	Y	Y	Y	Prob	Prob	RDI	Down	Lower Layer Down	Far End Alarm Rx
Y	Y	Y	Prob	Y	N/A	None	Up	Lower Layer Down	Incoming LOF
Y	Y	Y	Prob	Prob	N/A	RDI	Down	Lower Layer Down	Incoming LOF
Y	Prob	N/A	N/A	N/A	N/A	LOF	Down	Lower Layer Down	Incoming LOF
AIS	N/A	N/A	N/A	N/A	N/A	AIS	Down	Lower Layer Down	Incoming LOF
Prob	N/A	N/A	N/A	N/A	N/A	LOS	Down	Lower Layer Down	Incoming LOF

A DS3 path configured for PLCP mapping:

- Supports transmit and receive of the Ax, Px and C1 bits.
- Ignores the received Z1, Z2, Z3 octets of the PLCP frame and transmits all zeros in the Z1, Z2, Z3 octets of the PLCP frame.
- Ignores the received F1 octet of the PLCP frame, and transmits all zeros in the F1 octet of the PLCP frame.
- Samples and uses for performance monitoring received FEBE bits of G1 octet and transmits the number of BIP-8 errors detected by the receive framer using the FEBE bits of the G1 octet. Detects a PLCP Far End Alarm when 10 consecutive PLCP frames are received with the RAI bit set, and transmits a set RAI bit when the local port has declared PLCP-LOF. When the local port declares PLCP-LOF is cleared, the outgoing RAI bit is cleared.
- Ignores the received X bits of the G1 octet, and transmits all zeros in the X bits of the G1 octet of the PLCP frame.
- Ignores the received M1 and M2 octets and transmits all zeros in the M1 and M2 octets of the PLCP frame.

ATM Interface Configurations

Use the following CLI syntax to configure ATM interface parameters for SONET/SDH paths:

CLI Syntax:

```

config# port port-id
sonet-sdh
    path [sonet-sdh-index]
        atm
            cell-format cell-format
            ilmi [vpi/vci]
            egress
                traffic-desc traffic-desc-profile-id
            ingress
                traffic-desc traffic-desc-profile-id
            keep-alive [poll-frequency seconds] [poll-count value]
                [test-frequency seconds]
            protocol protocol-type
                [no] shutdown
            min-vp-vpi value

```

Use the following CLI syntax to configure ATM interface parameters for IMA bundles.

CLI Syntax:

```

config>port>multilink-bundle
ima
    atm
        cell-format cell-format
        min-vp-vpi value

```

Use the following CLI syntax to configure ATM interface parameters for TDM channels:

CLI Syntax:

```

config# port {port-id}
tdm
    ds1 [ds1-id]
        channel-group 1
            atm
                cell-format cell-format
                min-vp-vpi value
    ds3 [sonet-sdh-index]
        atm
            cell-format cell-format
            min-vp-vpi value
            mapping {direct | plcp}
    e1 [e1-id]

```

```

channel-group 1
    atm
        cell-format cell-format
        min-vp-vpi value
e3 [sonet-sdh-index]
    atm
        cell-format cell-format
        min-vp-vpi value

```

2.18.3.3.13 Configuring Frame Relay Parameters

Frame Relay pipes are used to provide customer-to-customer Frame Relay PVCs or to interconnect individual Frame Relay clouds.

Frame Relay parameters can only be configured in SONET/SDH and channelized TDM MDA contexts.

The following example shows a channelized interface configuration:

```

A:ALA-7>config>port# info detail
-----
description "DS3/E3"
...
    tdm
        buildout long
        ds3 ds3
            type t3
            channelized
            clock-source loop-timed
            framing c-bit
            no feac-loop-respond
            no mdl
            no mdl-transmit
            no loopback
            report-alarm ais los
            no report-alarm oof rai looped
            no shutdown
        exit
    ds1 ds1-1
        shutdown
        framing esf
        no loopback
        report-alarm ais los
        no report-alarm oof rai looped
        channel-group 1
            description "DS3/E3"
            mode access
            encap-type frame-relay
            no mtu
            no mac
            timeslots 1
            speed 64
            crc 16

```

```

        frame-relay
            lmi-type itu
            mode dte
            n393dce 4
            n393dte 4
            n391dte 6
            n392dce 3
            n392dte 3
            t391dte 10
            t392dce 15
        exit
    no shutdown
exit
exit
exit
no shutdown
-----
A:ALA-7>config>port#

```

SONET/SDH Interfaces

This section applies also to FR interfaces on Sonet/SDH high-speed channels on ASAP MDAs. In order to configure Frame Relay on the associated port/channel, the `frame-relay` encapsulation type must be specified.

The following output shows a Frame Relay encapsulation type and the Frame Relay defaults.

```

A:ALA-7>config>port# info detail
-----
    description "OC-3/OC-12 SONET/SDH"
    access
        ingress
            pool default
            resv-cbs default
            slope-policy "default"
        exit
    exit
    egress
        pool default
        resv-cbs sum
        slope-policy "default"
    exit
exit
network
    egress
        pool default
        resv-cbs default
        slope-policy "default"
    exit
exit
sonet-sdh
    framing sonet

```

```

clock-source node-timed
no loopback
speed oc12
report-alarm loc lrdi lb2er-sf slof slos
no report-alarm lais sslf lb2er-sd lrei
threshold ber-sd rate 6
threshold ber-sf rate 3
section-trace byte 0x1
path
    description "OC-3/OC-12 SONET/SDH"
    mode access
    encaps-type frame-relay
    no mtu
    no mac
    crc 32
    no scramble
    trace-string "Nokia 7750 ALA-"
    report-alarm plop pplm puneq
    no report-alarm pais prdi prei
    signal-label 0xcf
    frame-relay
        lmi-type itu
        mode dte
        n393dce 4
        n393dte 4
        n391dte 6
        n392dce 3
        n392dte 3
        t391dte 10
        t392dce 15
    exit
    no shutdown
exit
exit
no shutdown
-----
A:ALA-7>config>port# pwc

```

2.18.3.3.14 Configuring Multilink PPP Bundles

Multilink bundles can have from one to eight members (ports) specified. The bundles aggregate channelized ports which define available bandwidth to carry data over a DS1 channel. 56 multilink bundles can be configured per MDA. 256 MLPPP groups are supported per ASAP MDA. Each bundle represents a single connection between two routers.

Multilink bundling is based on a link control protocol (LCP) option negotiation that permits a system to indicate to its peer that it is capable of combining multiple physical links into a bundle.

Multilink bundling operations are modeled after a virtual PPP link-layer entity where packets received over different physical link-layer entities are identified as belonging to a separate PPP network protocol (the Multilink Protocol, or MP) and recombined and sequenced according to information present in a multilink fragmentation header. All packets received over links identified as belonging to the multilink arrangement are presented to the same network-layer protocol processing machine, whether or not they have multilink headers.

When you configure multilink bundles, consider the following guidelines:

- Multilink bundle configuration should include at least two ports.
- A maximum of eight ports can be included in a multilink bundle.
- Multilink bundles can only be aggregated on a single MDA.

```
A:ALA-A>config# port bundle-5/2.1
A:ALA-A>config>port# multilink-bundle
A:ALA-A>config>port>ml-bundle# member 5/2/1.ds0grp-1.1
A:ALA-A>config>port>ml-bundle# member 5/2/1.ds0grp-2.2
A:ALA-A>config>port>ml-bundle# member 5/2/1.ds0grp-1.1
```

2.18.3.3.15 Configuring Multilink ATM Inverse Multiplexing (IMA) Bundles

IMA bundles are supported on Channelized ASAP MDAs. The bundles aggregate E1 or DS1 ATM channels into a single logical ATM interface.

IMA Bundles

Use the following CLI syntax to configure IMA bundle parameters:

```
CLI Syntax: configure# port bundle-type-slot/mda.bundle-num
description description-string
multilink-bundle
    fragment-threshold value
    ima
        atm
            cell-format {uni|nni}
            min-vp-vpi vp-vpi-value
        exit
        link-delay {activate |deactivate} milli-seconds
        max-bandwidth number-links
        version ima-version
        red-differential-delay red-diff-delay down
        member port-id
```

Configuration notes:

An IMA group has common interface characteristics (for example, configuration that applies to a logical ATM interface either configured via the IMA group context or taken from the primary link) The following list details those common IMA group interface characteristics:

- Encapsulation type (ATM)
- ATM interface characteristics (under the ATM menu context)
- Interface mode type (only access is supported)
- MTU value (derived from the primary link)

Member links inherit those common characteristics from the IMA group that they are part of and as long as they are part of an IMA group. Characteristics derived from the primary link (MTU, interface mode type) can be changed on the primary link only and not on other links in the bundle or a bundle itself. The primary link is the member which has the lowest ifindex. When a member is added/deleted the primary member may be changed based on ifIndices of all member links.

Once a path becomes part of an IMA group logical link, the path ceases to exist as a physical ATM path interface. This means that:

1. ATM interface bundle characteristics enforced over the link. Note that when a link is removed from an IMA bundle, the link's ATM characteristics are reset to ATM interface defaults.
2. No services can be configured on the member link itself.

After the primary member has been added each additional member added to the group will only be accepted if it matches the configuration of the IMA group. ATM interface characteristics are not part of this verification as they are overwritten/reset to defaults when a link is added to/removed from an IMA bundle.

Upon addition to an IMA group, each added member is automatically assigned an IMA link ID. IMA link IDs are in range from 0 to 7 and stay constant as long as the router does not reboot.

When configuring IMA bundles, consider the following guidelines:

- IMA bundles should contain at least two members.
- A maximum of eight members can be included in an IMA bundle.
- IMA links can only be aggregated into a bundle within a single MDA.

- IMA group maximum bandwidth and minimum link settings allows, by default, for over-subscription of shaped services; however when that occurs scheduling of traffic over an IMA group ATM interface degrades to round-robin between shaped services, therefore to preserve full ATM TM even during a member link failure, it is recommended that maximum bandwidth is set to minimum links.
- When configuring the red differential delay for IMA groups on ASAP MDAs, the value configured is converted into acceptable frame sequence number delay on a link since delay is granular to IMA frame sequence number difference. For E1 channels (receiving frame time 27 ms), configured values map to the enforced values as follows: 0 ms maps to 0 frame sequence number difference (27 ms delay), 1 to 27 ms maps to 1 frame sequence number difference (54 ms delay), 28 - 50 ms maps to 2 frame sequence number difference (81 ms delay). Similarly, for DS1 channels (receiving frame time 35 ms), configured values map to enforced values as follows: 0 ms maps to 0 frame sequence number difference (35 ms delay), 1 to 35 ms maps to 1 frame sequence number difference (70 ms delay), 36 to 50 ms maps to 2 frame sequence number difference (105 ms delay).
- When a channel is deleted from an IMA group it is recommended that a deletion takes place at the far end first when the far end supports graceful deletion to ensure no cell loss takes place on the 7750 SR RX end of the channel. When a channel is deleted on the 7750 SR end first, a small data loss will take place on the 7750 SR RX side (depending on the time required for the far end to deactivate its TX on the link being deleted).
- When no member links are configured on an IMA group, the speed of an E1 channel will be used to compute the maximum IMA group bandwidth that may be allocated to shaped services.
- The shutdown command for IMA groups sets the IMA group state to "Blocking". This makes the group operationally down but will not bring down the individual IMA links. Services configured on the IMA group will go operationally down as well.
- The 7750 SR supports automatic IMA version changing when the far end IMA group version matches the configured version. The group will remain operationally down until one of the IMA groups changes version.
- When adding member links to an IMA group, the clock-source of the e1 or ds1 link must be set to node-timed.

The following example shows the creation of an IMA bundle with three group members residing on a channelized OC-3 ASAP MDA in slot 5/2/1:

```
A:ALA-A>config# port bundle-ima-5/2.1
A:ALA-A>config>port# multilink-bundle
A:ALA-A>config>port>ml-bundle# member 5/2/1.1.1.1
A:ALA-A>config>port>ml-bundle# member 5/2/1.1.2.1

A:ALA-A>config>port>ml-bundle# member 5/2/1.1.3.1
```

2.18.3.3.16 Multi-Class MLPPP

The following guidelines apply to multi-class MLPPP:

- MC-MLPPP must be configured before links are added to a bundle.
- MC-MLPPP and LFI (**config>port>ml-bundle>interleave-fragments**) are mutually exclusive.
- MC-MLPPP is not supported when port is configured as **network** mode.
- MC-MLPPP can be enabled on every MLPPP bundle and bundle protection group.
- MC-MLPPP is supported only on ASAP MDAs (for example, m4-choc3-as-sfp, m1-choc12-as-sfp, m4-chds3-as, m12-chds3-as).
- Short and long sequence packet formats are supported (both ends must be of the same type) with static mapping of forwarding classes to MC-MLPPP class (based on the number of classes negotiated with the far end).
- Single fragment size for all classes is supported.
- Prefix elision is not supported. The prefix elision (compressing common header bytes) option advises the peer that, in each of the given classes, the implementation expects to receive only packets with a certain prefix; this prefix is not to be sent as part of the information in the fragment(s) of this class.
- Fractional DS1/E1 MLPPP links are supported. This is applicable to MLPPP bundles on ASAP MDAs. Fractional E1 and Fractional DS1 links cannot be combined in the same bundle.

IMA Test Procedure

Use the following CLI commands to perform an IMA Test Pattern Procedure on a member link of an IMA group:

CLI Syntax:

```
configure# port bundle-type-slot/mda.bundle-num
multilink-bundle
    ima
    test-pattern-procedure
        test-link port-id
        test-pattern [pattern]
    no shutdown
```

An operator can deploy IMA test procedures to verify operations of IMA group and its member links. The following is a list of key points about the test pattern procedure:

1. The test procedure is performed as defined by the IMA specification version 1.1, i.e. a test pattern is sent over the specified link and is expected to be looped back over all the links in the group. ICP cells are used to perform the test.
2. The test procedure is not traffic affecting, for example, data traffic will not be affected by the ongoing test.
3. There can only be a single test executed per an IMA group at any given time
4. The IMA member link must exist in the specified group for the command to be accepted.
5. The test-pattern-procedure must be shutdown before a new test-link value or test pattern is accepted.
6. The current IMA group test pattern configuration and result of a given IMA test can be seen by executing a show command for the IMA group. A test-link result can have three values:
 - a. Disabled: The test-link is currently not running.
 - b. Operating: The test pattern procedure is **no shutdown** and there are currently no failed-links for this running test-pattern-procedure.
 - c. Link-Failed: One or more links have failed the test-pattern-procedure. Execute a **show port <slot/mda/port.sonet-sdh-index> ima-link** command to see the failed link and received pattern value.
7. Deleting a member link that is the same as the specified test-link, to stay in compliance with key point 4, will result in the test-link value being reset to default.
8. IMA test procedure configurations are not saved when the admin save command is executed.

2.18.3.3.17 Configuring Bundle Protection Group Ports

Bundle Protection groups enable APS protection of one bundle residing on a working circuit of an APS group port by another bundle residing on the protection circuit of that APS group port. Bundle protection groups apply to MLPPP as well, and are configured the same way. The following examples show the process to configure BPGp on ASAP MDAs to provide an APS protection for an IMA/MLPPP bundle.

First, two ASAP MDAs must be configured.

Example:

```
config# card 3
config>card# mda 2
config>card>mda# mda-type m4-choc3-as-sfp
config>card>mda# no shutdown
config>card>mda# exit
config>card# exit
config# card 10
```

```
config>card# mda 2
config>card>mda# mda-type m4-choc3-as-sfp
config>card>mda# no shutdown
config>card>mda# exit
```

Configure an APS group with working and protection circuits on the ASAP MDAs.

Example:

```
config# port aps-1
config>port# aps
config>port>aps# working-circuit 3/2/1
config>port>aps# protect-circuit 10/2/1
config>port>aps# exit
config>port# no shutdown
```

Create eight ATM DS1 channels on the APS group.

Example:

```
config>port>aps#
config>port# sonet-sdh
config>port>sonet-sdh# path sts1-1
config>port>sonet-sdh>path# no shutdown
config>port>sonet-sdh>path# exit
config>port>sonet-sdh# exit
config>port# tdm
config>port>tdm#
config>port>tdm# ds3 1
config>port>tdm>ds3# channelized ds1
config>port>tdm>ds3# no shutdown
config>port>tdm>ds3# exit
config>port>tdm# ds1 1.1
config>port>tdm>ds1# channel-group 1
config>port>tdm>ds1>channel-group# encap-type atm
config>port>tdm>ds1>channel-group# no shutdown
config>port>tdm>ds1>channel-group# exit
config>port>tdm# ds1 1.8
config>port>tdm>ds1# channel-group 1
config>port>tdm>ds1>channel-group# encap-type atm
config>port>tdm>ds1>channel-group# no shutdown
config>port>tdm>ds1>channel-group# exit
```

Next, configure an IMA-type/MLPPP-type BPGGrp with working and protection bundles on working and protection circuits of aps-1 and members the created DS1s (this creates 2 IMA bundles, one on working and one on protection circuit):

Example:

```
config# port bpgrp-ima-1
config>port# multilink-bundle
config>port>multilink-bundle# working-bundle bundle-ima-1/1.1
```

```
config>port>multilink-bundle# protect-bundle bundle-ima-
2/1.1
config>port>multilink-bundle# member aps-1.1.1.1
config>port>multilink-bundle# member aps-1.1.2.1
config>port>multilink-bundle# member aps-1.1.3.1
config>port>multilink-bundle# member aps-1.1.4.1
config>port>multilink-bundle# member aps-1.1.5.1
config>port>multilink-bundle# member aps-1.1.6.1
config>port>multilink-bundle# member aps-1.1.7.1
config>port>multilink-bundle# member aps-1.1.8.1
config>port>multilink-bundle# exit
config>port>multilink-bundle# no shutdown
config>port>multilink-bundle# exit
config>port# no shutdown
```

Finally, a service can be configured on this bundle using the BPGGrp ID (for example, an ATM VC 0/32 SAP would be: `sap bpg-ima-1:0/32`).

Configuration Notes and Guidelines:

- Any configuration on a BPGGrp applies to both the working and protection bundle.
- Working and protection bundles can be shutdown individually.
- Services cannot be configured on a BPGGrp until at least one member link has been configured.
- The published switchover times for bundle protection groups on the router are dependent on the far end being able to recover from cell loss within that time. To ensure this, the following recommendations are given:
 - The BPGGrp link activation timer should be configured to a value small enough to allow a quick recovery from any IMA failure occurring during the switchover. A recommended value is 1 second.
 - The ADM that terminates APS should support standard APS switchover time requirements.
 - The far end IMA/MLPPP links must be able to tolerate cell loss during APS switchover without bringing links down. This includes, for example, a combination of link activation/deactivation and appropriate configuration of TDM/SONET debounce timers.
 - Because of the temporary cell loss during the APS switchover, the far end IMA/MLPPP will experience a misalignment between individual links within an IMA/MLPPP group. The far end IMA/MLPPP group must support fast-realignment of links without having to bring the links down. The router synchronizes the IMA/MLPPP streams the far end receives between switchovers in an effort to cause the least amount of misalignment.

- To increase the BPGp robustness, it is recommended to provision more IMA/MLPPP links than is required and set the minimum links and max bandwidth parameters to the number of required links. This type of configuration is required on the far end as well.

Configuring a Channelized DS1 Card

The 7750 SR-c12 and 7750 SR-c4 support channelized DS-1 cards. The channelization is as follows:

- N*DS0 in DS1 port.{1 to 24}
- N*DS0 in E1 port.{1 to 32}

To ensure you have a channel-capable MDA or CMA, verify the MDA-type that you are configuring by entering a **show mda slot-id** command.

In the following example, MDA 7 shows a channelized DS1 CMA.

```
A:7710-3>config# show mda
=====
MDA Summary
=====
```

Slot	Mda	Provisioned Mda-type	Equipped Mda-type	Admin State	Operational State
1	1	m60-10/100eth-tx	m60-10/100eth-tx	up	up
	5	c8-10/100eth-tx	c8-10/100eth-tx	up	up
	6	c1-lgb-sfp	c1-lgb-sfp	up	up
	7	c8-chds1	c8-chds1	up	up
	8	c4-ds3	c4-ds3	up	up

```
=====
A:7710-3>

A:7710-3>config# show mda 1/7 detail
=====
MDA 1/7 detail
=====
```

Slot	Mda	Provisioned Mda-type	Equipped Mda-type	Admin State	Operational State
	7	c8-chds1	c8-chds1	up	up

```
=====
MDA Specific Data
Maximum port count      : 8
Number of ports equipped : 8
Network ingress queue policy : default
Capabilities            : TDM, PPP, FR
Min channel size        : PDH DS0 Group
Max channel size         : PDH DS1
Max number of channels   : 64
Channels in use          : 0

Hardware Data
```

```

Part number           : Sim Part#
CLEI code             : Sim CLEI
Serial number         : mda-7
Manufacture date      : 01012003
Manufacturing string   : Sim MfgString mda-7
Manufacturing deviations : Sim MfgDeviation mda-7
Administrative state   : up
Operational state     : up
Temperature           : 35C
Temperature threshold  : 75C
Time of last boot      : 2006/10/02 09:28:22
Current alarm state    : alarm cleared
Base MAC address      : 04:7b:01:07:00:01

```

```

=====
A:7710-3>

```

In the TDM E1 context, configure DS0 channel groups and their parameters. For a DS1 channel-group, up to 24 timeslots can be assigned (numbered 1 to 24). For an E1 channel-group, up to 31 timeslots can be assigned (numbered 2 to 32). For ATM, all timeslots are auto-configured when a channel group gets created (there is no sub-E1 for ATM). ATM, Frame Relay and BCP-NUL encapsulation examples follow:

```

ALA-A>config>port>tdm# e1 1.1
ALA-A>config>port>tdm>e1# channel-group 1
ALA-A>config>port>tdm>e1>channel-group# timeslots 2
ALA-A>config>port>tdm>e1>channel-group# no shutdown
ALA-A>config>port>tdm>e1>channel-group#
ALA-A>config>port>tdm>e1# no shutdown
ALA-A>config>port>tdm>e1# channel-group 2
ALA-A>config>port>tdm>e1>channel-group# timeslots 3
ALA-A>config>port>tdm>e1>channel-group# encap-type frame-relay
ALA-A>config>port>tdm>e1>channel-group# no shutdown
ALA-A>config>port>tdm>e1>channel-group# exit
ALA-A>config>port>tdm>e1# channel-group 3
ALA-A>config>port>tdm>e1>channel-group# timeslots 11,12
ALA-A>config>port>tdm>e1>channel-group# encap-type cisco-hdlc
ALA-A>config>port>tdm>e1>channel-group# no shutdown
ALA-A>config>port>tdm>e1>channel-group# exit
ALA-A>config>port>tdm>e1# no shutdown
ALA-A>config>port>tdm>e1# exit
ALA-A>config>port>tdm# e1 1.2
ALA-A>config>port>tdm>e1# no shutdown
ALA-A>config>port>tdm>e1# channel-group 1
ALA-A>config>port>tdm>e1>channel-group# encap-type atm
ALA-A>config>port>tdm>e1>channel-group# no shutdown
ALA-A>config>port>tdm>e1>channel-group# exit
ALA-A>config>port>tdm>e1# no shutdown
ALA-A>config>port>tdm# info
-----
tdm
  ds3 1
    no shutdown
  exit
  e1 1.1
    channel-group 1
      timeslots 2
    no shutdown

```



```

        exit
        channel-group 2
        encap-type frame-relay
        frame-relay
        exit
        timeslots 10
        no shutdown
    exit
    channel-group 3
    encap-type cisco-hdlc
    cisco-hdlc
    exit
    timeslots 11,12
    no shutdown
exit
no shutdown
exit
e1 1.2
    channel-group 1
    encap-type atm
    atm
    exit
    no shutdown
exit
no shutdown
exit
no shutdown
-----
ALA-A>config>port>tdm#

```

Services can now be applied to the configured channelized ports.

2.18.3.3.18 Configuring LAG Parameters

LAG configurations should include at least two ports. Other considerations include:

- A maximum of 64 ports (depending on the lag-id) can be included in a LAG. All ports in the LAG must share the port characteristics inherited from the primary port.
- Autonegotiation must be disabled or set limited mode for ports that are part of a LAG to guarantee a specific port speed.
- Ports in a LAG must be configured as full duplex.

The following example shows LAG configuration output:

```

A:ALA-A>config>lag# info detail
-----
description "LAG2"
mac 04:68:ff:00:00:01
port 1/1/1
port 1/3/1
port 1/5/1

```

```
port 1/7/1
port 1/9/1
dynamic-cost
port-threshold 4 action down
-----
A:ALA-A>config>lag#
```

Configuring BFD on LAG Links

BFD can be configured under the LAG context to create and establish the micro-BFD session per link after the LAG and associated links have been configured. An IP interface must be associated with the LAG or a VLAN within the LAG, if dot1q encapsulation is used, before the micro-BFD sessions can be established.

Complete the following steps to enable and configure BFD over the individual LAG links:

- Step 1.** Enable BFD within the LAG context, which also enters the CLI into the BFD context
- Step 2.** Configure the address family which is to be used for the micro BFD sessions. Only one address family can be configured per LAG
- Step 3.** Configured the local-IP address to be used for the BFD sessions
- Step 4.** Configure the remote-IP address to be used for the BFD sessions

When configuring the local and remote IP address for the BFD over LAG link sessions, the *local-ip* parameter should always match an IP address associated with the IP interface to which this LAG is bound. In addition, the *remote-ip* parameter should match an IP address on the remote system and should also be in the same subnet as the *local-ip* address. If the LAG bundle is re-associated with a different IP interface, the *local-ip* and *remote-ip* parameters should be modified to match the new IP subnet. The *local-ip* and *remote-ip* values do not have to match a configured interface in the case of tagged LAG/ports.

The optional parameters that can be configured for the BFD over LAG links include:

- Transmit Interval
- Receive Interval
- Multiplier
- Max-Wait-for-Up-Time - This parameter controls how long a link will remain active if BFD is enabled after the LAG and associated links are active and in a forwarding state.
- Max-Time-Admin-Down - This parameter controls how long the system will wait before bringing the associated link out of service if an admin down message is received from the far-end.

The following is a sample configuration:

```
*A:Dut-C>config>lag# info
-----
bfd
  family ipv4
    local-ip-address 10.120.1.2
    receive-interval 1000
    remote-ip-address 10.120.1.1
    transmit-interval 1000
    no shutdown
  exit
exit
no shutdown
```

2.18.3.3.19 Configuring G.8031 Protected Ethernet Tunnels

Ethernet tunnel configuration can include at most two paths. Other considerations include:

- A path contains one member port and one control-tag (backbone VLAN ID/ BVID)
- If the operator wants to replace an existing member port or a control-tag, the whole path needs to be shutdown first. The alternate path will be activated as a result keeping the traffic interruption to a minimum. Then the whole path must be deleted and re-created. To replace an existing member port or control tag, the whole path needs to be shutdown first. The alternate path will be activated as a result keeping traffic interruption to a minimum. Then the whole path must be deleted, the alternate path precedence modified to primary before re-creating the new path.
- The Ethernet tunnel will inherit the configuration from the first member port. The following port-level configuration needs to be the same between member ports of an Ethernet tunnel:
 - config>port>ethernet>access>{ingress|egress}>queue-group
 - config>port>ethernet>egress-scheduler-policy
 - config>port>access>egress>pool
 - config>port>ethernet>dot1q-etype
 - config>port>ethernet>qinq-etype
 - config>port>ethernet>pbb-etype
 - config>port>ethernet> mtu
- The operator can update these port parameters only if the port is the sole member of an Ethernet tunnel. This means that in the example below, the operator needs to remove port 1/1/4 and port 1/1/5 before being allowed to modify 1/1/1 for the above parameters.

CLI Syntax:

```
eth-tunnel 1
  path 1
    member 1/1/1
  path 2
    member 1/1/4
eth-tunnel 2
  path 1
    member 1/1/1
  path 2
    member 1/1/5
```

The following example shows eth-tunnel configuration output:

```
port 1/1/1
  ethernet
    encap-type dot1q
port 2/2/2
  ethernet
    encap-type dot1q

config eth-tunnel 1
  path 1
    member 1/1/1
    control-tag 100
    precedence primary
    eth-cfm
      mep 51 domain 1 association 1
      ccm-enable
      low-priority-defect allDef
      mac-address 00:AE:AE:AE:AE:AE
      control-mep
      no shutdown
  no shutdown
  path 2
    member 2/2/2
    control-tag 200
eth-cfm
  mep
    mep 52 domain 1 association 2 direction down
    ccm-enable
    low-priority-defect allDef
    mac-address 00:BE:BE:BE:BE:BE
    control-mep
    no shutdown
no shutdown
```

2.18.3.3.20 Configuring Connectors and Connector Ports

Some assemblies have support for QSFP28 or QSFP-DD transceiver modules. These modules have different variants, some of which provide multiple physical ports out of a single module (breakout modules). There is a QSFP28 breakout module that supports four physical 10 Gb Ethernet ports. On assemblies that support these breakout variants, the front panel cages are modeled as connectors rather than as direct ports. The connector must be configured for the type of breakout module that is to be inserted and then the appropriate ports are created and can be configured.

The connector reference is in the format *slot/mda/cconnector* (for example, 1/1/c3) and the ports owned by the connector use the format *slot/mda/cconnector/port*. For example, in a 7750 SR-1 with the 6-port QSFP28 mda-e-xp installed in the first MDA slot, initially there are no ports available, only six connectors:

```
A:bkvm18# show mda

=====
MDA Summary
=====
Slot  Mda  Provisioned Type                               Admin   Operational
                               Equipped Type (if different)   State   State
-----
1      1      me6-100gb-qsfp28                               up      up
*A:bkvm18# show port

=====
Ports on Slot 1
=====
Port      Admin Link Port   Cfg  Oper  LAG/  Port  Port  Port  C/QS/S/XFP/
Id         State      State  MTU  MTU   Bndl  Mode  Encp  Type  MDIMDX
-----
1/1/c1     Up      No    Down                -  unkn  unkn  conn
1/1/c2     Up      No    Down                -  unkn  unkn  conn
1/1/c3     Up      No    Down                -  unkn  unkn  conn
1/1/c4     Up      No    Down                -  unkn  unkn  conn
1/1/c5     Up      No    Down                -  unkn  unkn  conn
1/1/c6     Up      No    Down                -  unkn  unkn  conn
=====
```

After configuring a module with four 10 Gb breakout ports in connector position 1/1/c1 and a module with one 100 Gb breakout port in connector position 1/1/c2, the physical ports are created:

```
A:bkvm18# configure port 1/1/c1 connector breakout c4-10g
A:bkvm18# configure port 1/1/c2 connector breakout c1-100g

A:bkvm18>show port

=====
Ports on Slot 1
=====
Port      Admin Link Port   Cfg  Oper  LAG/  Port  Port  Port  C/QS/S/XFP/
Id         State      State  MTU  MTU   Bndl  Mode  Encp  Type  MDIMDX
-----
```

1/1/c1	Up	No	Link Up			- unkn unkn conn	40GBASE-SR4
1/1/c1/1	Down	No	Down	9212	9212	- netw null xgige	
1/1/c1/2	Down	No	Down	9212	9212	- netw null xgige	
1/1/c1/3	Down	No	Down	9212	9212	- netw null xgige	
1/1/c1/4	Down	No	Down	9212	9212	- netw null xgige	
1/1/c2	Down	No	Link Up			- unkn unkn conn	100GBASE-LR4
1/1/c2/1	Down	No	Down	1578	1578	- netw null cgige	
1/1/c3	Up	No	Down			- unkn unkn conn	
1/1/c4	Up	No	Down			- unkn unkn conn	
1/1/c5	Up	No	Down			- unkn unkn conn	
1/1/c6	Up	No	Down			- unkn unkn conn	

These physical ports can now be used as Ethernet port references in other commands.

The transceiver information is shown under the connector while Ethernet-related items are shown under the connector ports.

*A:Dut-A# show port 1/1/c1

```
=====
QSFP28 Connector
=====
Description      : QSFP28 Connector
Interface        : 1/1/c1
Admin State      : up
Oper State       : up
IfIndex          : 104939520
Last State Change : 10/31/2017 13:23:22
Last Cleared Time  : N/A
Breakout         : c4-10g
DMM Events       : Enabled

Transceiver Data

Transceiver Status : operational
Transceiver Type    : QSFP28
Model Number       : 3HE10551AARA01 NOK IPU3BFVEAA
TX Laser Wavelength: 850 nm
Number of Lanes     : 4
Connector Code      : MPO 1x12
Vendor OUI         : 44:7c:7f
Manufacture date    : 2017/03/12
Media               : Ethernet
Serial Number       : INHAG8480890
Part Number        : TR-FC85S-NNO
Optical Compliance : 100GBASE-SR4 or 25GBASE-SR
Link Length support: 70m for OM3; 100m for OM4

=====
Transceiver Digital Diagnostic Monitoring (DDM)
=====
Value High Alarm High Warn Low Warn Low Alarm
-----
Temperature (C)      +29.8    +80.0    +75.0    -5.0    -10.0
Supply Voltage (V)    3.29     3.63     3.46     3.14     2.97
=====

=====
Transceiver Lane Digital Diagnostic Monitoring (DDM)
=====
```

	High Alarm	High Warn	Low Warn	Low Alarm
Lane Tx Bias Current (mA)	12.0	10.0	4.5	3.0
Lane Tx Output Power (dBm)	5.40	2.40	-8.40	-11.40
Lane Rx Optical Pwr (avg dBm)	5.40	2.40	-10.30	-13.30

Lane ID	Temp (C) /Alm	Tx Bias (mA) /Alm	Tx Pwr (dBm) /Alm	Rx Pwr (dBm) /Alm
1	-	7.1	-0.08	-4.56
2	-	0.0/L-WA	-40.00/L-WA	-40.00/L-WA
3	-	0.0/L-WA	-40.00/L-WA	-40.00/L-WA
4	-	0.0/L-WA	-40.00/L-WA	-40.00/L-WA

```

=====
*A:Dut-A#
*A:Dut-A#
*A:Dut-A# show port 1/1/c1/1
=====

```

Ethernet Interface

```

=====
Description      : 10-Gig Ethernet
Interface        : 1/1/c1/1
Link-level       : Ethernet
Admin State      : up
Oper State       : up
Physical Link     : Yes
Single Fiber Mode : No
IfIndex          : 104939521
Last State Change : 10/31/2017 13:23:23
Last Cleared Time  : N/A
Phys State Chng Cnt: 1

Configured Mode   : network
Dot1Q Ethertype   : 0x8100
PBB Ethertype     : 0x88e7
Ing. Pool % Rate  : 100
Ing. Pool Policy  : n/a
Egr. Pool Policy  : n/a
Net. Egr. Queue Pol: default
Egr. Sched. Pol   : n/a
HS Scheduler Plcy : default
HS Port Pool Plcy : default
Monitor Port Sched : Disabled
Monitor Agg Q Stats: Disabled
Auto-negotiate    : N/A
Oper Phy-tx-clock : not-applicable
Accounting Policy : None
Acct Plcy Eth Phys : None
Egress Rate       : Default
Load-balance-algo : Default
Access Bandwidth  : Not-Applicable
Access Available BW: 0
Access Booked BW  : 0
Sflow             : Disabled

Oper Speed        : 10 Gbps
Config Speed      : N/A
Oper Duplex       : full
Config Duplex     : N/A
MTU               : 9212
Min Frame Length  : 64 Bytes
Hold time up      : 0 seconds
Hold time down    : 0 seconds

RS-FEC Mode       : None

Encap Type        : null
QinQ Ethertype    : 0x8100

Egr. Pool % Rate  : 100

MDI/MDX           : N/A

Collect-stats     : Disabled
Collect Eth Phys  : Disabled
Ingress Rate      : Default
LACP Tunnel       : Disabled
Booking Factor    : 100

Suppress Threshold : 2000
Reuse Threshold    : 1000
=====

```

```

Max Penalties      : 16000                      Max Suppress Time: 20 seconds
Half Life          : 5 seconds

Down-when-looped   : Disabled                    Keep-alive       : 10
Loop Detected      : False                       Retry           : 120
Use Broadcast Addr : False

Sync. Status Msg.  : Disabled                    Rx Quality Level : N/A
Tx DUS/DNU         : Disabled                    Tx Quality Level : N/A
SSM Code Type      : sdh

Down On Int. Error : Disabled                    DOIE Tx Disable  : N/A

CRC Mon SD Thresh  : Disabled                    CRC Mon Window   : 10 seconds
CRC Mon SF Thresh  : Disabled

EFM OAM            : Disabled                    EFM OAM Link Mon : Disabled
Ignr EFM OAM State : False

Configured Address : 00:03:fa:2b:ef:1a
Hardware Address   : 00:03:fa:2b:ef:1a
Cfg Alarm          : remote local

```

```

=====
Traffic Statistics
=====

```

	Input	Output
Octets	0	0
Packets	0	0
Errors	0	0
Utilization (300 seconds)	0.00%	0.00%

```

=====
Port Statistics
=====

```

	Input	Output
Unicast Packets	0	0
Multicast Packets	0	0
Broadcast Packets	0	0
Discards	0	0
Unknown Proto Discards	0	

```

=====
Ethernet-like Medium Statistics
=====

```

Alignment Errors :	0	Sngl Collisions :	0
FCS Errors :	0	Mult Collisions :	0
SQE Test Errors :	0	Late Collisions :	0
CSE :	0	Excess Collisns :	0
Too long Frames :	0	Int MAC Tx Errs :	0
Symbol Errors :	0	Int MAC Rx Errs :	0
In Pause Frames :	0	Out Pause Frames :	0

2.19 Service Management Tasks

This section discusses basic procedures to complete service management tasks.

2.19.1 Modifying or Deleting an MDA, MCM, CMA or XMA

To change an MDA, MCM, CMA, or XMA type already provisioned for a specific slot or card, first you must shut down the slot/MDA/port configuration and then delete the MDA, MCM, CMA, or the XMA from the configuration.

To modify or delete CMAs or XMA, use the MDA command structure.

Use the following CLI syntax to modify an MDA on the 7450 ESS and 7750 SR platforms (or an XMA on the 7950 XRS platforms):

CLI Syntax: `config> port port-id
shutdown`

CLI Syntax: `config> card slot-number
shutdown
[no] mda mda-number
[no] mda-type mda-type
shutdown`



Note: You do not have to shutdown and remove an MCM to remove or modify an MDA on the 7750 SR. Use the following sequence if changing the MCM type or slot configuration.

2.19.2 Modifying a Card Type

In order to modify the card type already provisioned for a specific slot, you must shutdown existing port configurations and shutdown and remove all MDA, XMA, or CMA configurations. For 7750 SR-c12/c4 systems, after removing MDA configurations, shutdown and remove the MCM from service before modifying the card.

CMAs do not require an MCM, therefore, if removing a CMA-type MDA from service, it is not required to shutdown and remove an MCM before modifying the card.

You must reset the IOM after changing the MDA type from MS-ISA to any other MDA type.

Use the following CLI syntax to modify a card type already provisioned for a specific slot:

CLI Syntax: `config> port port-id`
 `[no] shutdown`

CLI Syntax: `config> card slot-number`
 `mda mda-number`
 `[no] mda-type mda-type`
 `[no] shutdown`

CLI Syntax: `config> card slot-number`
 `shutdown`
 `[no] mcm mcm-number`
 `no mcm-type mcm-type`
 `shutdown`

2.19.3 Deleting a Card

To delete a card type provisioned for a specific slot, you must shutdown existing port configurations and shutdown and remove all MDA, XMA, or CMA configurations. For 7750 SR-c12/c4 systems, after removing MDA configurations, you can shutdown and remove the MCM from service before modifying the card.

Use the following CLI syntax to delete a card provisioned for a specific slot:

CLI Syntax: `config> port port-id`
 `shutdown`

CLI Syntax: `config> card slot-number`
 `card-type card-type`
 `mda mda-number`
 `no mda-type mda-type`
 `no shutdown`
 `mcm mcm-number (for 7750 SR-c12/c4 only)`
 `no mcm-type mcm-type`
 `no shutdown`

2.19.4 Deleting Port Parameters

Use the following CLI syntax to delete a port provisioned for a specific card or CMA:

CLI Syntax: `config>port port-id`
 `shutdown`
 `no port port-id`

2.19.5 Soft IOM Reset

This section provides basic procedures for soft IOM reset service management tasks.

2.19.5.1 Soft Reset

Soft reset is an advanced high availability feature that greatly reduces the impact of IOM/IMM resets either during a software upgrade or during other maintenance or debug operations. The combination of In Service Software Upgrade (ISSU) and Soft reset maximizes service availability in an operational network.

A soft reset re-initializes the control plane while the data plane continues operation with only very minimal impact to data forwarding. During the soft reset some processes that rely on the IOM control plane will not run for a duration that is similar to the duration of an IOM Hard reset. These processes include the updating of the IP forwarding table on the IOM (IP FIB downloads from the CPM), Layer 2 learning of new MAC addresses on the IOM, updating of the MAC forwarding table (for MAC addresses learned from other IOMs), ARP, Ethernet OAM 802.3ah, LLDP and handling for certain ICMP functions such as Can't Fragment, Redirect, Host Unreachable, Network Unreachable and TTL Expired. Note that protocols and processes on the CPM continue to operate during a Soft Reset (BGP continues to learn new routes from peers, and the new routes will be downloaded to the IOM once the Soft Reset has completed).

The combination of the very small data plane impact and special soft reset enhancements for protocols ensures that most protocols do not go down and no visible impacts to most protocols are detected externally to the SR/ESS platforms. BFD timers are temporarily increased for the duration of a soft reset in order to keep BFD sessions up. Protocols such as BGP, OSPF, IS-IS, PIM, etc with default timers remain up. A protocol using aggressive timers may go down momentarily during a soft reset.

Although the majority of protocols stay up during a Soft Reset, there are some limitations for a few protocols. Refer to *Known Limitations* in the *Release Notes* for the relevant release for details.

Configuration changes are not allowed while any card is in the process of a soft reset.

The soft IOM reset procedure is applicable during the ISSU process and for a manual soft reset procedure.

To manually perform a soft IOM reset, enter the **clear card slot-number soft** command.

Soft Reset is supported on Ethernet IMMs and on IOMs that have Ethernet MDAs provisioned. The operator can optionally force a Soft Reset on an IOM that contains at least one MDA that supports Soft Reset but also has an MDA that does not support Soft Reset or is operationally down. To force Soft Reset in this case the **hard-reset-unsupported-mdas** is used and the supported MDAs and the card itself are soft reset while the MDAs that do not support soft reset (or are operationally down) are hard reset.

The **show card** and **show mda** commands indicate that a soft IOM reset is occurring during the soft reset process.

Soft Reset is not supported on the 7750 SR-c4. On the 7750 SR-c12, Soft Reset is not supported but the ISSU procedure will avoid resetting soft reset capable MDAs/CMAs (as long as there is not new firmware for the CMA/MDA in the new image).

2.19.5.2 Deferred MDA Reset

As part of an ISSU, soft reset is supported even if the (old) firmware version on the MDAs is not the same as the (new) firmware version in the software load to which the operator is upgrading. The soft reset is allowed to proceed by leaving the previous version of the firmware running while upgrading the rest of the MDA/IOM/IMM. The operator can then issue a hard reset of the MDA/IMM at some time in the future to upgrade the firmware.

The soft reset is only allowed to proceed if the older firmware is compatible with the new IOM/IMM software load. Otherwise the soft reset is blocked and a hard reset must be used instead.

After a soft reset has been completed, a log event will be raised to warn the operator that the MDA (or IMM) is running older firmware and that they can perform a hard reset of the MDA (or IMM) at some point if required.

If the MDA/IMM is not hard reset by the operator, and then a software upgrade is performed, and the older firmware is no longer compatible with the newest load being upgraded to, then the soft reset will be blocked (or an automatic hard reset will occur for Major ISSU).

The operator can see whether they are running with older MDA/IMM firmware at any time by using the `show mda detail` command.

2.20 Configuration Command Reference

2.20.1 Command Hierarchies

- [Card Commands](#)
- [MCM Commands](#)
- [MDA Commands](#)
- [Power Commands](#)
- [Virtual Scheduler Commands](#)
- [Forwarding Plane \(FP\) Commands](#)
- [Port Configuration Commands](#)
- [Port XC Commands](#)
- [Forwarding Path Extension \(FPE\) Commands](#)
- [Port APS Commands](#)
- [Ethernet Commands](#)
- [Ethernet Access and Network Commands](#)
- [HSQ Port Configuration Commands](#)
- [Interface Group Handler Commands](#)
- [Multilink Bundle Commands](#)
- [SONET-SDH Commands](#)
- [TDM Commands](#)
- [DS3 Commands](#)
- [E1 Commands](#)
- [E3 Commands](#)
- [LAG Commands](#)
- [MACsec Commands](#)
- [Ethernet Tunnel Commands](#)
- [Multi-Chassis Redundancy Commands](#)

2.20.1.1 Card Commands

```
config
  — [no] card slot-number
    — card-type card-type [level card-level]
```

- no **card-type**
- [no] **fail-on-error**
- [no] **named-pool-mode** [now]
- [no] **reset-on-recoverable-error**
- [no] **shutdown**
- **upgrade** *index path upgrade-name*
- no **upgrade** *index*

2.20.1.2 MCM Commands

- ```

config
— [no] card slot-number
— [no] mcm mcm-slot
— mcm-type mcm-type
— no mcm-type
— [no] shutdown

```

### 2.20.1.3 MDA Commands

- ```

config
— [no] card slot-number
— [no] mda mda-slot
— access
— egress
— [no] pool [name]
— amber-alarm-threshold percentage
— no amber-alarm-threshold
— red-alarm-threshold percentage
— no red-alarm-threshold
— resv-cbs percent-or-default amber-alarm-action step
— percent max percent
— resv-cbs percent-or-default
— no resv-cbs
— slope-policy name
— no slope-policy
— ingress
— [no] pool [name]
— amber-alarm-threshold percentage
— no amber-alarm-threshold
— red-alarm-threshold percentage
— no red-alarm-threshold
— resv-cbs percent-or-default amber-alarm-action step
— percent max percent
— resv-cbs percent-or-default
— no resv-cbs
— slope-policy name
— no slope-policy
— clock-mode adaptive

```



```

— clock-mode differential [timestamp-freq {19440 | 77760 | 103680}]
— egress
— egress-xpl
    — threshold xpl-errors
    — window minutes
— [no] fail-on-error
— ingress-xpl
    — threshold threshold
    — window window
— mda-type mda-type [level mda-level]
— no mda-type
— named-pool-mode
    — egress
        — named-pool-policy policy-name
        — no named-pool-policy
    — ingress
        — named-pool-policy policy-name
        — no named-pool-policy
— network
    — egress
        — [no] pool [name]
            — amber-alarm-threshold percentage
            — no amber-alarm-threshold
            — red-alarm-threshold percentage
            — no red-alarm-threshold
            — resv-cbs percent-or-default amber-alarm-action step
                percent max percent
            — resv-cbs percent-or-default
            — no resv-cbs
            — slope-policy name
            — no slope-policy
        — ingress
            — [no] pool [name]
                — amber-alarm-threshold percentage
                — no amber-alarm-threshold
                — red-alarm-threshold percentage
                — no red-alarm-threshold
                — resv-cbs percent-or-default amber-alarm-action step
                    percent max percent
                — resv-cbs percent-or-default
                — no resv-cbs
                — slope-policy name
                — no slope-policy
            — queue-policy name
            — no queue-policy
    — power-priority-level priority
    — no power-priority-level
    — [no] reset-on-recoverable-error
    — [no] shutdown
    — [no] sync-e
    — upgrade index path upgrade-name
    — no upgrade index
— [no] shutdown

```

2.20.1.4 Power Commands

```

config
  — system
    — power-management
      — mode [none | basic | advanced]
      — [no] pcm pcm-slot [chassis chassis-id]
        — pcm-type peq-type
        — no pcm-type
      — [no] peq peq-slot [chassis chassis-id]
        — input-power-mode amperage
        — peq-type peq-type
        — no peq-type
        — [no] shutdown
      — power-safety-alert wattage
      — power-safety-level percent

```

2.20.1.5 Virtual Scheduler Commands

```

config
  — [no] card slot-number
    — virtual-scheduler-adjustment
      — internal-scheduler-weight-mode {default | force-equal | offered-load | capped-offered-load}
      — no internal-scheduler-weight-mode
      — rate-calc-min-int [fast-queue percent-of-default] [slow-queue percent-of-default]
      — no rate-calc-min-int
      — sched-run-min-int percent-of-default
      — no sched-run-min-int
      — slow-queue-threshold kilobits-per-second
      — no slow-queue-threshold
      — task-scheduling-int percent-of-default
      — no task-scheduling-int

```

2.20.1.6 Forwarding Plane (FP) Commands

```

config
  — [no] card slot-number
    — fp [fp-number]
      — egress
        — hs-fixed-high-thresh-delta size-in-bytes
        — no hs-fixed-high-thresh-delta
        — hs-pool-policy name
        — no hs-pool-policy
        — wred-queue-control
          — buffer-allocation min percentage max percentage

```

- **no buffer-allocation**
- **resv-cbs** *min percentage max percentage*
- **no resv-cbs**
- **[no] shutdown**
- **slope-policy** *slope-policy-name*
- **no slope-policy**
- **hi-bw-mcast-src** [*alarm*] [*group group-id*] [*default-paths-only*]
- **no hi-bw-mcast-src**
- **ingress**
 - **access**
 - **queue-group** *queue-group-name instance instance-id* [*create*]
 - **no queue-group** *queue-group-name instance instance-id*
 - **accounting-policy** *acct-policy-id*
 - **no accounting-policy**
 - **[no] collect-stats**
 - **description** *description-string*
 - **no description**
 - **policer-control-override** [*create*]
 - **no policer-control-override**
 - **max-rate** *{rate | max}*
 - **priority-mbs-thresholds**
 - **min-thresh-separation** *size [bytes | kilobytes]*
 - **[no] priority** *level*
 - **mbs-contribution** *[bytes | kilobytes]*
 - **policer-control-policy** *policer-control-policy-name*
 - **no policer-control-policy**
 - **[no] policer-override**
 - **policer** *policer-id* [*create*]
 - **no policer** *policer-id*
 - **cbs** *{size [bytes | kilobytes] | default}*
 - **no cbs**
 - **mbs** *{size [bytes | kilobytes] | default}*
 - **no mbs**
 - **packet-byte-offset** *{add add-bytes | subtract sub-bytes}*
 - **no packet-byte-offset**
 - **rate** *{rate | max} [cir {max | rate}]*
 - **no rate**
 - **stat-mode** *stat-mode*
 - **no stat-mode**
 - **mcast-path-management**
 - **bandwidth-policy** *policy-name*
 - **no bandwidth-policy**
 - **[no] shutdown**
 - **network**
 - **queue-group** *queue-group-name instance instance-id* [*create*]
 - **no queue-group** *queue-group-name instance instance-id*
 - **accounting-policy** *acct-policy-id*
 - **no accounting-policy**
 - **[no] collect-stats**
 - **description** *description-string*
 - **no description**

- **policer-control-override** *[create]*
- **no policer-control-override**
 - **max-rate** *{rate | max}*
 - **priority-mbs-thresholds**
 - **min-thresh-separation** *size [bytes | kilobytes]*
 - *[no] priority level*
 - **mbs-contribution** *size [bytes | kilobytes]*
- **policer-control-policy** *policer-control-policy-name*
- **no policer-control-policy**
- *[no] policer-override*
 - **policer** *policer-id [create]*
 - **no policer** *policer-id*
 - **cbs** *{size [bytes | kilobytes] | default}*
 - **no cbs**
 - **mbs** *{size [bytes | kilobytes] | default}*
 - **no mbs**
 - **packet-byte-offset** *{add bytes | subtract bytes}*
 - **no packet-byte-offset**
 - **rate** *{rate | max} [cir {max | rate}]*
 - **no rate**
 - **stat-mode** *stat-mode*
 - **no stat-mode**
- **ingress-buffer-allocation** *percentage*
- **no ingress-buffer-allocation**
- **policy-accounting** *limit*
- **no policy-accounting**
- *[no] stable-pool-sizing*

2.20.1.7 Port Configuration Commands

- ```
config
— [no] port {port-id | bundle-id | bpggrp-id | aps-id | connector-port-id}
 — access
 — egress
 — [no] pool [name]
 — amber-alarm-threshold percentage
```

- **no amber-alarm-threshold**
- **red-alarm-threshold** *percentage*
- **no red-alarm-threshold**
- **resv-cbs** *percent-or-default* **amber-alarm-action** **step** *percent*  
*max percent*
- **resv-cbs** *percent-or-default*
- **no resv-cbs**
- **slope-policy** *name*
- **no slope-policy**
- **connector**
  - **breakout** *breakout*
  - **no breakout**
- **[no] ddm-events**
- **description** *long-description-string*
- **no description**
- **dwdm**
  - **amplifier**
    - **[no] report-alarms** *[ild] [tmp] [mth] [mtl] [los] [lop] [com]*
  - **channel** *channel*
  - **coherent**
    - **channel** *channel*
    - **compatibility** *mode*
    - **cpr-window-size** *window-size*
    - **dispersion** *dispersion*
    - **mode** {**automatic** | **manual**}
    - **[no] report-alarms** *[modflt] [mod] [netrx] [nettx] [hosttx]*
    - **rx-los-reaction** {**squelch**}
    - **no rx-los-reaction**
    - **rx-los-thresh** *threshold*
    - **sweep** **start** *dispersion-start* **end** *dispersion-end*
    - **target-power** *power*
  - **[no] rxdtv-adjust**
  - **tdcm**
    - **channel** *channel*
    - **dispersion** *dispersion*
    - **mode** {**automatic** | **manual**}
    - **[no] report-alarms** *[nrly] [mth] [mtl] [unlck] [tlim] [einv] [com]*
    - **sweep** **start** *dispersion-start* **end** *dispersion-end*
  - **wavetracker**
    - **encode** **key1** *wave-key* **key2** *wave-key*
    - **no encode**
    - **[no] power-control**
      - **target-power** *dBm*
    - **[no] report-alarm** *[encode-fail] [encode-degrade] [power-fail]*  
*[power-degrade] [power-high] [power-low] [missing]*
- **hybrid-buffer-allocation**
  - **egr-weight** **access** *access-weight* **network** *network-weight*
  - **no egr-weight**
  - **ing-weight** **access** *access-weight* **network** *network-weight*
  - **no ing-weight**
- **modify-buffer-allocation-rate**
  - **egr-percentage-of-rate** *egr-rate-percentage*
  - **no egr-percentage-of-rate**
  - **ing-percentage-of-rate** *ing-rate-percentage*

- 
- **no ing-percentage-of-rate**
  - **[no] monitor-agg-egress-queue-stats**
  - **named-pool-mode**
    - **egress**
      - **named-pool-policy** *policy-name*
      - **no named-pool-policy**
    - **ingress**
      - **named-pool-policy** *policy-name*
      - **no named-pool-policy**
  - **network**
    - **egress**
      - **[no] pool** [*name*]
        - **amber-alarm-threshold** *percentage*
        - **no amber-alarm-threshold**
        - **red-alarm-threshold** *percentage*
        - **no red-alarm-threshold**
        - **resv-cbs** *percent-or-default* **amber-alarm-action** **step** *percent*  
max *percent*
        - **resv-cbs** *percent-or-default*
        - **no resv-cbs**
        - **slope-policy** *name*
        - **no slope-policy**
  - **[no] otu**
    - **[no] async-mapping**
    - **fec** {**g709** | **enhanced**}
    - **no fec**
    - **otu2-lan-data-rate** {**11.049** | **11.096**}
    - **pm-tti**
      - **expected** **auto-generated**
      - **expected** **bytes** *byte-string* [*byte-string* ... (up to 64 byte-strings max, 64 bytes max)]
      - **expected** **string** *identifier*
      - **expected** **use-rx**
      - **mismatch-reaction** {**squelch-rx**}
      - **no mismatch-reaction**
      - **tx** **auto-generated**
      - **tx** **bytes** *byte-string* [*byte-string* ... (up to 64 byte-strings max, 64 bytes max)]
      - **tx** **string** *identifier*
      - **no tx**
    - **psi-payload**
      - **expected** **byte**
      - **expected** **auto**
      - **mismatch-reaction** {**squelch-rx**}
      - **no mismatch-reaction**
      - **tx** **byte**
      - **tx** **auto**
    - **[no] psi-tti**
      - **expected** **auto-generated**
      - **expected** **bytes** *byte-string* [*byte-string* ... (up to 64 byte-strings max, 255 bytes max)]
      - **expected** **string** *identifier*
      - **expected** **use-rx**
      - **mismatch-reaction** {**squelch-rx**}

- **no mismatch-reaction**
- **tx** auto-generated
- **tx bytes** *byte-string* [*byte-string*...(up to 64 byte-strings max, 255 bytes max)]
- **tx string** *identifier*
- **no tx**
- [no] **report-alarms** [loc] [los] [lof] [lom] [otu-ais] [otu-ber-sf] [otu-ber-sd] [otu-bdi] [otu-tim] [otu-iae] [otu-biae] [fec-sf] [fec-sd] [fec-fail] [fec-uncorr] [odu-ais] [odu-oci] [odu-lck] [odu-bdi] [odu-tim] [opu-tim] [opu-plm]
- **sf-sd-method** {bip8 | fec}
- **sf-threshold** *threshold* [*coefficient* *coefficient*]
- **sf-threshold-clear** *threshold* [*coefficient* *coefficient*]
- **sm-tti**
  - **expected** auto-generated
  - **expected bytes** *byte-string* [*byte-string*...(up to 64 byte-strings max, 64 bytes max)]
  - **expected string** *identifier*
  - **expected use-rx**
  - **mismatch-reaction** {*squelch-rx*}
  - **no mismatch-reaction**
  - **tx** auto-generated
  - **tx bytes** *byte-string* [*byte-string*...(up to 64 byte-strings max, 64 bytes max)]
  - **tx string** *identifier*
  - **no tx**
- [no] **shutdown**

### 2.20.1.8 Port XC Commands

- ```
config
— port-xc
  — pxc pxc-id [create]
  — no pxc pxc-id
    — description description-string
    — no description
    — port port-id
    — no port
    — [no] shutdown
— [no] port
  — pxc-pxc-id.sub-port-id
```

2.20.1.9 Forwarding Path Extension (FPE) Commands

- ```
config
— fwd-path-ext
 — fpe fpe-id [create]
 — no fpe fpe-id
```

- **description** *description-string*
- **no description**
- **path** **xc-a** *lag-id* **xc-b** *lag-id*
- **path** **pxc** *pxc-id*
- **no path**
- **[no] pw-port**
- **[no] sub-mgmt-extensions**
- **vxlan-termination** [**router** *router-name*]
- **vxlan-termination** **service-name** *service-name*
- **no vxlan-termination**
- **sdp-id-range** **from** *id* **to** *id*
- **no sdp-id-range**

### 2.20.1.10 Port APS Commands

- ```
config
— [no] port {aps-id}
  — aps
    — advertise-interval advertise-interval- no advertise-interval
- hold-time hold-time
- no hold-time
- hold-time-aps [!signal-failure sf-time] [!signal-degrade sd-time]
- no hold-time-aps
- [no] mode-annexb
- neighbor ip-address
- no neighbor
- protect-circuit port-id
- no protect-circuit
- rdi-alarms [suppress | circuit]
- revert-time minutes
- no revert-time
- switching-mode {bi-directional | uni-directional | uni-1plus1}
- working-circuit port-id [number number]
- no working-circuit [number number]
- wtr-annexb minutes

```

2.20.1.11 Ethernet Commands

- ```
config
— [no] port {port-id}
 — ethernet
 — accounting-policy policy-id- no accounting-policy
- autonegotiate [limited]
- no autonegotiate
- [no] collect-stats
- crc-monitor

```



- **sd-threshold** *threshold* [**multiplier** *multiplier*]
- **no sd-threshold**
- **sf-threshold** *threshold* [**multiplier** *multiplier*]
- **no sf-threshold**
- **window-size** *seconds*
- **no window-size**
- **dampening**
  - **half-life** *half-life* **max-suppress-time** *max-time*
  - **[no] shutdown**
  - **suppress-threshold** *suppress-penalties* **reuse-threshold** *reuse-penalties*
- **dot1q-etype** *value*
- **no dot1q-etype**
- **dot1x**
  - **[no] macsec**
    - **exclude-mac-policy** *mac-policy-id*
    - **no exclude-mac-policy**
    - **[no] exclude-protocol** *{protocol-name}*
    - **[no] rx-must-be-encrypted**
    - **sub-port** *port-id* **[create]**
    - **no sub-port** *port-id*
      - **ca-name** *ca-name*
      - **no ca-name**
      - **eapol-destination-address** *mac*
      - **no eapol-destination-address**
      - **encap-match** *{all-encap | double-tag encap-value | single-tag encap-value | untagged}*
      - **no encap-match**
      - **max-peer** *max-peer*
      - **no max-peer**
      - **[no] shutdown**
  - **max-auth-req** *max-auth-request*
  - **port-control** *{auto | force-auth | force-unauth}*
  - **quiet-period** *seconds*
  - **radius-plcy** *name*
  - **no radius-plcy**
  - **re-auth-period** *seconds*
  - **no re-auth-period**
  - **[no] re-authentication**
  - **server-timeout** *seconds*
  - **no server-timeout**
  - **supplicant-timeout** *seconds*
  - **no supplicant-timeout**
  - **transmit-period** *seconds*
  - **no transmit-period**
  - **no tunneling**
- **down-on-internal-error** **[tx-disable]**
- **no down-on-internal-error**
- **down-when-looped**
  - **keep-alive** *timer*
  - **no keep-alive**
  - **retry-timeout** *timer*
  - **no retry-timeout**
  - **[no] shutdown**

- 
- [no] **use-broadcast-address**
  - **duplex** {full | half}
  - **efm-oam**
    - [no] **accept-remote-loopback**
    - **discovery**
      - **advertise-capabilities**
        - [no] **link-monitoring**
    - [no] **dying-gasp-tx-on-reset**
    - [no] **grace-tx-enable**
    - **grace-vendor-oui** *oui*
    - **no grace-vendor-oui**
    - **hold-time** *time-value*
    - **no hold-time**
    - [no] **ignore-efm-state**
    - **link-monitoring**
      - **errored-frame**
        - [no] **event-notification**
        - **sd-threshold** *errored-frames*
        - **no sd-threshold**
        - **sf-threshold** *errored-frames*
        - [no] **shutdown**
        - **window** *deciseconds*
      - **errored-frame-period**
        - [no] **event-notification**
        - **sd-threshold** *errored-frames*
        - **no sd-threshold**
        - **sf-threshold** *errored-frames*
        - [no] **shutdown**
        - **window** *packets*
      - **errored-frame-seconds**
        - [no] **event-notification**
        - **sd-threshold** *errored-seconds*
        - **no sd-threshold**
        - **sf-threshold** *errored-seconds*
        - [no] **shutdown**
        - **window** *deciseconds*
      - **errored-symbols**
        - [no] **event-notification**
        - **sd-threshold** *errored-symbols*
        - **no sd-threshold**
        - **sf-threshold** *errored-symbols*
        - [no] **shutdown**
        - **window** *deciseconds*
      - **local-sf-action**
        - **event-notification-burst** *packets*
        - **info-notification**
          - [no] **critical-event**
          - [no] **dying-gasp**
        - **local-port-action** {log-only | out-of-service}
      - [no] **shutdown**
    - **mode** {active | passive}
    - **peer-rdi-rx**
      - **critical-event** **local-port-action** {log-only | out-of-service}
      - **dying-gasp** **local-port-action** {log-only | out-of-service}

- **event-notification** local-port-action {log-only | out-of-service}
- **link-fault** local-port-action {log-only | out-of-service}
- [no] **shutdown**
- [no] **transmit-interval** interval [multiplier multiplier]
- **trigger-fault** {dying-gasp | critical-event}
- **no trigger-fault**
- [no] **tunneling**
- **egress**
  - [no] **exp-secondary-shaper** secondary-shaper-name [create]
    - **agg-burst**
      - **high-burst-increase** size [bytes | kilobytes]
      - **no high-burst-increase**
      - **low-burst-limit** size [bytes | kilobytes]
      - **no low-burst-limit**
    - **class** class-number rate rate [monitor-threshold size-in-kbytes] [burst-limit size] [bytes | kilobytes]
    - **no class** class-number
    - **low-burst-max-class** class
    - **no low-burst-max-class**
    - **rate** rate [monitor-threshold size-in-bytes]
    - **no rate**
  - **hs-port-pool-policy** policy-name
  - **no hs-port-pool-policy**
  - **hs-scheduler-overrides** [create]
  - **no hs-scheduler-overrides**
    - **group** group-id rate rate
    - **no group** group-id
    - **max-rate** rate
    - **no max-rate**
    - **scheduling-class** class rate rate
    - **scheduling-class** class weight weight-in-group
    - **no scheduling-class** class
  - **hs-scheduler-policy** policy-name
  - **no hs-scheduler-policy**
  - **hs-secondary-shaper** secondary-shaper-name [create]
  - **no hs-secondary-shaper** secondary-shaper-name
    - [no] **aggregate**
      - **low-burst-max-class** class
      - **no low-burst-max-class**
      - **rate** rate
      - **no rate**
    - [no] **class** class-number
      - **rate** rate
      - **no rate**
    - **description** description-string
    - **no description**
- **egress-rate** sub-rate
- **no egress-rate**
- **egress-scheduler-override** [create]
- **no egress-scheduler-override**
  - **level** priority-level rate pir-rate [cir cir-rate]
  - **level** priority-level percent-rate pir-percent [percent-cir cir-percent]
  - **no level** priority-level

---

```

— max-rate rate
— max-rate percent percent-rate
— no max-rate
— egress-scheduler-policy port-scheduler-policy-name
— no egress-scheduler-policy
— elmi
 — mode {none | uni-n}
 — n393 [value]
 — no n393
 — t391 [value]
 — no t391
 — t392 [value]
 — no t392
— encap-type {dot1q | null | qinq}
— no encap-type
— eth-bn-egress-rate-changes
— no eth-bn-egress-rate-changes
— eth-cfm
 — [no] mep mep-id domain md-index association ma-index [vlan vlan-id]
 — [no] ais-enable
 — client-meg-level [level [level]]
 — no client-meg-level
 — [no] interface-support-enable
 — interval {1 | 60}
 — no interval
 — low-priority-defect {allDef | macRemErrXcon}
 — priority priority-value
 — no priority
 — alarm-notification
 — fng-alarm-time time
 — fng-reset-time time
 — [no] ccm-enable
 — ccm-ltm-priority priority
 — no ccm-ltm-priority
 — ccm-padding-size ccm-padding
 — no ccm-padding-size
 — ccm-tlv-ignore [port-status] [interface-status]
 — no ccm-tlv-ignore
 — [no] collect-lmm-stats
 — [no] csf-enable
 — multiplier multiplier-value
 — no multiplier
 — description description-string
 — no description
 — eth-bn
 — [no] receive
 — rx-update-pacing seconds
 — [no] eth-test-enable
 — bit-error-threshold bit-errors
 — test-pattern {all-zeros | all-ones} [crc-enable]
 — no test-pattern
 — [no] facility-fault
 — grace
 — eth-ed

```

---

```

— max-rx-defect-window seconds
— no max-rx-defect-window
— priority priority
— no priority
— [no] rx-eth-ed
— [no] tx-eth-ed
— eth-vsm-grace
— [no] rx-eth-vsm-grace
— [no] tx-eth-vsm-grace
— low-priority-defect {allDef | macRemErrXcon | remErrXcon |
errXcon | xcon | noXcon}
— mac-address mac-address
— no mac-address
— one-way-delay-threshold seconds
— [no] shutdown
— hold-time {[up hold-time-up] [down hold-time-down] [seconds |
centiseconds]}
— no hold-time
— [no] hsm-da-scheduler-overrides
— group group-id rate rate
— no group group-id
— max-rate rate
— no max-rate
— scheduling-class class rate rate
— scheduling-class class weight weight-in-group
— no scheduling-class class
— ingress-rate ingress-rate
— no ingress-rate
— [no] lacp-tunnel
— lldp
— dest-mac {nearest-bridge | nearest-non-tpmr | nearest-customer}
— admin-status {rx | tx | tx-rx | disabled}
— [no] notification
— port-id-subtype {tx-if-alias | tx-if-name | tx-local}
— [no] tunnel-nearest-bridge
— tx-mgmt-address [system] [system-ipv6] [oob] [oob-ipv6]
— no tx-mgmt-address
— tx-tlvs [port-desc] [sys-name] [sys-desc] [sys-cap]
— no tx-tlvs
— load-balancing-algorithm option
— no load-balancing-algorithm
— mac ieee-address
— no mac
— min-frame-length byte-length
— mode {access | network | hybrid}
— no mode
— [no] mon-port-sch
— mtu mtu-bytes
— no mtu
— pbb-etype [ethertype-value]
— no pbb-etype
— ptp-asymmetry nanoseconds
— no ptp-asymmetry
— qinq-etype qinq-etype-value

```

- **no qinq-etype**
- **[no] report-alarm** *[signal-fail] [remote] [local] [no-frame-lock] [high-ber] [no-block-lock] [no-am-lock] [duplicate-lane]*
- **rs-fec-mode** *rs-fec-mode*
- **no rs-fec-mode**
- **[no] sflow**
- **[no] single-fiber**
- **speed** {10 | 100 | 1000 | 10000 | 25000 | 40000 | 50000 | 100000}
- **ssm**
  - **code-type** {sonet | sdh}
  - **no code-type**
  - **[no] shutdown**
  - **[no] tx-dus**
- **symbol-monitor**
  - **sd-threshold** *threshold [multiplier multiplier]*
  - **no sd-threshold**
  - **sf-threshold** *threshold [multiplier multiplier]*
  - **no sf-threshold**
  - **[no] shutdown**
  - **window-size** *seconds*
  - **no window-size**
- **util-stats-interval** *seconds*
- **xgig** {lan | wan}

```

config
 — system
 — ethernet
 — efm-oam
 — [no] dying-gasp-tx-on-reset

```

### 2.20.1.11.1 Ethernet Access and Network Commands

```

config>port>ethernet
 — access
 — bandwidth bandwidth
 — no bandwidth
 — booking-factor factor
 — no booking-factor
 — egress
 — [no] queue-group queue-group-name [create] [instance instance-id]
 — accounting-policy acct-policy-id
 — no accounting-policy
 — [no] agg-rate
 — [no] limit-unused-bandwidth
 — [no] queue-frame-based-accounting
 — rate kilobits-per-second
 — no rate
 — [no] collect-stats

```

- **description** *description-string*
- **no description**
- **host-match** *destination-string* [**create**]
- **no host-match** *destination-string*
- **[no] hs-turbo**
- **[no] hsmda-queue-override**
  - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
  - **no packet-byte-offset**
  - **queue** *queue-id* [**create**]
  - **no queue**
    - **mbs** {[**0** to **2625**] **kilobytes** | [**0** to **2688000**] **bytes** | **default**}
    - **no mbs**
    - **rate** *pir-rate*
    - **no rate**
    - **slope-policy** *hsmda-slope-policy-name*
    - **no slope-policy**
    - **wrr-weight** *weight*
    - **no wrr-weight**
  - **secondary-shaper** *secondary-shaper-name*
  - **no secondary-shaper**
  - **wrr-policy** *wrr-policy-name*
- **queue-overrides**
  - **queue** *queue-id* [**create**]
  - **no queue** *queue-id*
    - **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
    - **no adaptation-rule**
    - **burst-limit** {**default** | *size* [**bytes** | **kilobytes**]}
    - **no burst-limit**
    - **cbs** *size-in-kbytes*
    - **no cbs**
    - **drop-tail**
      - **low**
        - **percent-reduction-from-mbs** *percent*
        - **no percent-reduction-from-mbs**
    - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
    - **no mbs**
    - **[no] monitor-depth**
    - **parent** [*weight weight*] [*cir-weight cir-weight*]
    - **no parent**
    - **percent-rate** *pir-percent* [*cir cir-percent*]
    - **no percent-rate**
    - **rate** *pir-rate* [*cir cir-rate*]
    - **no rate**
- **[no] scheduler-override**
  - **scheduler** *scheduler-name* [**create**]
  - **no scheduler** *scheduler-name*
    - **parent** {[*weight weight*] [*cir-weight cir-weight*]}
    - **no parent**
    - **rate** *pir-rate* [*cir cir-rate*]
    - **no rate**
- **scheduler-policy** *scheduler-policy-name*
- **no scheduler-policy**

---

```

— vport name [create]
— no vport name
 — [no] agg-rate
 — [no] limit-unused-bandwidth
 — rate kilobits-per-second
 — no rate
 — description description-string
 — no description
 — [no] egress-rate-modify
 — host-match dest description-string [create]
 — no host-match dest destination-string
 — [no] mon-port-sch
 — port-scheduler-policy port-scheduler-policy-name
 — no port-scheduler-policy
 — scheduler-policy scheduler-policy-name
 — no scheduler-policy
— ingress
 — [no] queue-group queue-group-name [create]
 — accounting-policy acct-policy-id
 — no accounting-policy
 — [no] collect-stats
 — description description-string
 — no description
 — queue-overrides
 — queue queue-id [create]
 — no queue queue-id
 — adaptation-rule [pir adaptation-rule] [cir {max | min |
 closest}]
 — no adaptation-rule
 — cbs size-in-kbytes
 — no cbs
 — drop-tail
 — low
 — percent-reduction-from-mbs percent
 — no percent-reduction-from-mbs
 — mbs {size [bytes | kilobytes] | default}
 — no mbs
 — [no] monitor-depth
 — rate pir-rate [cir cir-rate]
 — no rate
 — [no] scheduler-override
 — scheduler scheduler-name [create]
 — no scheduler scheduler-name
 — parent {[weight weight] [cir-weight cir-weight]}
 — no parent
 — rate pir-rate [cir cir-rate]
 — no rate
 — scheduler-policy scheduler-policy-name
 — no scheduler-policy
— network
 — accounting-policy policy-id
 — no accounting-policy
 — [no] collect-stats
 — egress

```



- 
- [no] **queue-group** *queue-group-name* [**instance** *instance id*] [**create**]
    - **accounting-policy** *acct-policy-id*
    - **no accounting-policy**
    - **agg-rate**
    - **no agg-rate**
      - [no] **limit-unused-bandwidth**
      - [no] **queue-frame-based-accounting**
      - **rate** *kilobits-per-second*
      - **no rate**
    - [no] **collect-stats**
    - **description** *description-string*
    - **no description**
    - [no] **hs-turbo**
    - **policer-control-policy** *policy-name*
    - **queue-overrides**
      - **queue** *queue-id* [**create**]
      - **no queue** *queue-id*
        - **adaptation-rule** [**pir** *adaptation-rule*] [**cir** {**max** | **min** | **closest**}]
        - **no adaptation-rule**
        - **cbs** *size-in-kbytes*
        - **no cbs**
        - **drop-tail**
          - **low**
            - **percent-reduction-from-mbs** *percent*
            - **no percent-reduction-from-mbs**
        - **mbs** {*size* [**bytes** | **kilobytes**] | **default**}
        - **no mbs**
        - [no] **monitor-depth**
        - **percent-rate** *pir-percent* [**cir** *cir-percent*]
        - **no percent-rate**
        - **rate** *pir-rate* [**cir** *cir-rate*]
        - **no rate**
      - **parent** *scheduler-policy-name*
      - **no parent**
      - **scheduler-policy** *scheduler-policy-name*
      - **no scheduler-policy**
    - **queue-policy** *name*
    - **no queue-policy**

## 2.20.1.12 HSQ Port Configuration Commands

```

config>port>ethernet
 — access
 — egress
 — [no] queue-group queue-group-name [create] [instance instance-id]
 — [no] hs-turbo
 — egress
 — hs-port-pool-policy policy-name
 — no hs-hs-port-pool-policy
 — hs-scheduler-overrides [create]
 — no hs-scheduler-overrides
 — group group-id rate rate
 — no group group-id
 — max-rate rate
 — no max-rate
 — scheduling-class class rate rate
 — scheduling-class class weight weight-in-group
 — no scheduling-class class
 — hs-scheduler-policy policy-name
 — no hs-scheduler-policy
 — hs-secondary-shaper secondary-shaper-name [create]
 — no hs-secondary-shaper secondary-shaper-name
 — [no] aggregate
 — low-burst-max-class class
 — no low-burst-max-class
 — rate rate
 — no rate
 — [no] class class-number
 — rate rate
 — no rate
 — description description-string
 — no description
 — network
 — egress
 — [no] queue-group queue-group-name [create] [instance instance-id]
 — [no] hs-turbo

```

## 2.20.1.13 Interface Group Handler Commands

```

config
 — [no] interface-group-handler index
 — [no] member port-id
 — [no] shutdown
 — threshold num-members
 — no threshold

```

## 2.20.1.14 Multilink Bundle Commands

```

config
 — [no] port {bundle-id}
 — multilink-bundle
 — fragment-threshold fragment-threshold
 — fragment-threshold unlimited
 — no fragment-threshold
 — ima
 — atm
 — cell-format {uni | nni}
 — min-vp-vpi value
 — link-delay {activate | deactivate} milli-seconds
 — no link-delay {activate | deactivate}
 — max-bandwidth number-links
 — no max-bandwidth
 — test-pattern-procedure
 — [no] shutdown
 — test-link port-id
 — no test-link
 — test-pattern pattern
 — no test-pattern
 — version IMA-version
 — no version
 — [no] interleave-fragments
 — [no] member port-id
 — minimum-links minimum-links
 — no minimum-links
 — mlfr
 — ack-timeout timer
 — no ack-timeout
 — egress
 — qos-profile profile-id
 — no qos-profile
 — frame-relay
 — lmi-type {ansi | itu | none | rev1}
 — mode {dce | dte | bidir}
 — n391dte intervals
 — no n391dte
 — n392dce threshold
 — no n392dce
 — n392dte threshold
 — no n392dte
 — n393dce count
 — no n393dce
 — n393dte count
 — no n393dte
 — t391dte keepalive
 — no t391dte
 — t392dce keepalive
 — no t392dce
 — hello-interval timer
 — no hello-interval

```

---

```

— identifier frf16-identifier
— no identifier
— ingress
 — qos-profile profile-id
 — no qos-profile
— retry-limit count
— no retry-limit
— mlppp
 — egress
 — qos-profile profile-id
 — no qos-profile
 — endpoint-discriminator class {ip-address | global-mac-address
 |null} [discriminator-id discriminator-id]
 — no endpoint-discriminator
 — ingress
 — qos-profile profile-id
 — no qos-profile
 — [no] magic-number
 — multiclass count
 — no multiclass
 — [no] stateless-aps-switchover
— mrru mrru
— no mrru
— [no] protect-bundle bundle-id
— red-differential-delay red-diff-delay [down]
— no red-differential-delay
— [no] short-sequence
— [no] working-bundle bundle-id
— yellow-differential-delay yellow-diff-delay
— no yellow-differential-delay

```

### 2.20.1.15 SONET-SDH Commands

```

config
— [no] port {port-id}
 — sonet-sdh
 — clock-source {loop-timed | node-timed}
 — framing {sonet | sdh}
 — group sonet-sdh-index payload {tu3 | vt2 | vt15}
 — hold-time {[up hold-time-up] [down hold-time-down] }
 — no hold-time
 — loopback {line | internal}
 — no loopback
 — [no] path [sonet-sdh-index]
 — access
 — egress
 — vport name [create]
 — no vport name
 — [no] agg-rate
 — [no] limit-unused-bandwidth
 — [no] queue-frame-based-accounting

```

- **rate** *kilobits-per-second*
  - **no rate**
  - **description** *description-string*
  - **no description**
  - **[no] egress-rate-modify**
  - **host-match dest** *description-string* [**create**]
  - **no host-match dest** *destination-string*
  - **[no] mon-port-sch**
  - **port-scheduler-policy** *port-scheduler-policy-name*
  - **no port-scheduler-policy**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
- **atm**
  - **cell-format** {**uni** | **nni**}
  - **[no] custom-buffer-mode**
    - **buffer-pool** *value*
    - **vc-threshold** *buffer-threshold*
  - **ilmi** [*vpi/vci*]
  - **no ilmi**
    - **egress**
      - **traffic-desc** *traffic-desc-profile-id*
      - **no traffic-desc**
    - **ingress**
      - **traffic-desc** *traffic-desc-profile-id*
      - **no traffic-desc**
    - **keep-alive** [*poll-frequency seconds*] [*poll-count value*] [*test-frequency seconds*]
    - **no keep-alive**
    - **protocol** *protocol-type*
    - **no protocol**
    - **[no] shutdown**
    - **min-vp-vpi** *value*
  - **cisco-hdlc**
    - **down-count** *down-count*
    - **no down-count**
    - **keepalive** *time-interval*
    - **no keepalive**
    - **up-count** *up-count*
    - **no up-count**
  - **crc** {**16** | **32**}
  - **description** *long-description-string*
  - **no description**
  - **egress-scheduler-override** [**create**]
  - **no egress-scheduler-override**
    - **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*]
    - **level** *priority-level* **percent-rate** *pir-percent* [**percent-cir** *cir-percent*]
    - **no level** *priority-level*
    - **max-rate** *rate*
    - **max-rate** **percent** *percent-rate*
    - **no max-rate**
  - **egress-scheduler-policy** *port-scheduler-policy-name*
  - **no egress-scheduler-policy**

- 
- **encap-type** {atm | bcp-null | bcp-dot1q | ipcp | ppp-auto | frame-relay | wan-mirror | cisco-hdlc}
  - **frame-relay**
    - [no] **frf-12**
      - **egress**
        - **qos-profile** *profile-id*
        - **no qos-profile**
      - **fragment-threshold** *threshold*
      - **no fragment-threshold**
    - **lmi-type** {ansi | itu | none | rev1}
    - **mode** {dce | dte | bidir}
    - **n391dte** *intervals*
    - **no n391dte**
    - **n392dce** *threshold*
    - **no n392dce**
    - **n392dte** *threshold*
    - **no n392dte**
    - **n393dce** *count*
    - **no n393dce**
    - **n393dte** *count*
    - **no n393dte**
    - **t391dte** *keepalive*
    - **no t391dte**
    - **t392dce** *keepalive*
    - **no t392dce**
  - **load-balancing-algorithm** *option*
  - **no load-balancing-algorithm**
  - **mac** *ieee-address*
  - **no mac**
  - **mode** {access | network | hybrid}
  - **mtu** *mtu-bytes*
  - **no mtu**
  - **network**
    - **accounting-policy** *policy-id*
    - **no accounting-policy**
    - [no] **collect-stats**
    - **queue-policy** *name*
    - **no queue-policy**
  - **payload** {sts3 | tug3 | ds3 | e3 | vt2 | vt15 | ds1 | e1}
  - **ppp**
    - **keepalive** *time-interval* [**dropcount** *drop-count*]
    - **no keepalive**
  - [no] **report-alarm** [pais] [plop] [prdi] [pplm] [prei] [puneq] [plcd]
  - [no] **scramble**
  - [no] **shutdown**
  - **signal-label** *value*
  - **no signal-label**
  - **trace-string** [*trace-string*]
  - **no trace-string**
- [no] **report-alarm** [loc] [lais] [lrldi] [ss1f] [lb2er-sd] [lb2er-sf] [slof] [slos] [lrei]
  - [no] **reset-port-on-path-down**
  - **section-trace** {increment-z0 | **byte** *value* | **string** *string*}
  - [no] **single-fiber**
  - **speed** {oc3 | oc12}

- **no speed**
- **[no] suppress-lo-alarm**
- **threshold** {ber-sd | ber-sf} rate *threshold-rate*
- **no threshold** {ber-sd | ber-sf}
- **[no] tx-dus**

## 2.20.1.16 TDM Commands

```

config
 — [no] port {port-id}
 — tdm
 — buildout {long | short}
 — [no] ds1 ds1-id
 — bert {2e3 | 2e9 | 2e11 | 2e15 | 2e20 | 2e20q | 2e23 | ones | zeros |
 alternating} duration duration
 — no bert
 — bit-error-insertion rate
 — no bit-error-insertion
 — [no] channel-group channel-group
 — atm
 — cell-format {uni | nni}
 — min-vp-vpi value
 — cisco-hdlc
 — down-count down-count
 — no down-count
 — keepalive time-interval
 — no keepalive
 — up-count up-count
 — no up-count
 — crc {16 | 32}
 — description long-description-string
 — no description
 — egress-scheduler-override [create]
 — no egress-scheduler-override
 — level priority-level rate pir-rate [cir cir-rate]
 — level priority-level percent-rate pir-percent [percent-cir
 cir-percent]
 — no level priority-level
 — max-rate rate
 — max-rate percent percent-rate
 — no max-rate
 — egress-scheduler-policy port-scheduler-policy-name
 — encap-type {atm | bcp-null | bcp-dot1q | ipcp | ppp-auto |
 frame-relay | wan-mirror | cisco-hdlc | cem}
 — frame-relay
 — [no] frf-12
 — egress
 — qos-profile profile-id
 — no qos-profile
 — fragment-threshold threshold
 — no fragment-threshold

```

---

```

— identifier frf16-identifier
— no identifier
— lmi-type {ansi | itu | none | rev1}
— mode {dce | dte | bidir}
— n391dte intervals
— no n391dte
— n392dce threshold
— no n392dce
— n392dte threshold
— no n392dte
— n393dce count
— no n393dce
— n393dte count
— no n393dte
— t391dte keepalive
— no t391dte
— t392dce keepalive
— no t392dce
— idle-cycle-flag {flags | ones}
— no idle-cycle-flag
— idle-payload-fill {all-ones}
— idle-payload-fill pattern pattern
— no idle-payload-fill
— idle-signal-fill {all-ones}
— idle-signal-fill pattern pattern
— no idle-signal-fill
— load-balancing-algorithm option
— no load-balancing-algorithm
— mac ieee-address
— no mac
— mode {access | network}
— mtu mtu-bytes
— no mtu
— network
— accounting-policy policy-id
— no accounting-policy
— [no] collect-stats
— queue-policy name
— no queue-policy
— ppp
— [no] ber-sf-link-down
— compress {acfc [pfc] | pfc [acfc]}
— no compress
— keepalive time-period [dropcount drop count]
— no keepalive
— [no] scramble
— [no] shutdown
— speed {56 | 64}
— timeslots timeslots
— no timeslots
— clock-source {loop-timed | node-timed | adaptive | differential}
— framing (DS-1) {esf | sf | ds1-unframed}
— [no] invert-data

```



- **loopback** {line | internal | fdl-ansi | fdl-bellcore | payload-ansi | inband-ansi | inband-bellcore}
- **no loopback**
- [no] **remote-loop-respond**
- [no] **report-alarm** [ais] [los] [oof] [rai] [looped] [ber-sd] [ber-sf]
- [no] **shutdown**
- **signal-mode** {cas}
- **no signal-mode**
- **threshold** {ber-sd | ber-sf} rate {1 | 5 | 10 | 50 | 100}
- **no threshold** {ber-sd | ber-sf}
- **hold-time** {[up *hold-time-up*] [down *hold-time-down*]}
- **no hold-time**
- **lbo** [0dB | -7.5dB | -15.0dB | -22.5dB]
- **length** {133 | 266 | 399 | 533 | 655}

## 2.20.1.17 DS3 Commands

- ```
config
  — [no] port {port-id}
    — tdm
      — [no] ds3 [sonet-sdh-index]
        — atm
          — cell-format {uni | nni}
          — mapping mapping
          — min-vp-vpi value
        — bert {2e3 | 2e9 | 2e11 | 2e15 | 2e20 | 2e20q | 2e23 | ones | zeros | alternating} duration duration
        — no bert
        — bit-error-insertion rate
        — no bit-error-insertion
        — channelized {ds1 | e1}
        — no channelized
        — cisco-hdlc
          — down-count down-count
          — no down-count
          — keepalive time-interval
          — no keepalive
          — up-count up-count
          — no up-count
        — clock-source {loop-timed | node-timed}
        — crc {16 | 32}
        — description long-description-string
        — no description
        — egress-scheduler-override [create]
        — no egress-scheduler-override
          — level priority-level rate pir-rate [cir cir-rate]
          — level priority-level percent-rate pir-percent [percent-cir cir-percent]
          — no level priority-level
          — max-rate rate
          — max-rate percent percent-rate
```

- **no max-rate**
- **egress-scheduler-policy** *port-scheduler-policy-name*
- **no egress-scheduler-policy**
- **encap-type** {atm | bcp-null | bcp-dot1q | ipcp | ppp-auto | frame-relay | wan-mirror | cisco-hdlc | cem}
- [no] **feac-loop-respond**
- **frame-relay**
 - [no] **frf-12**
 - **egress**
 - **qos-profile** *profile-id*
 - **no qos-profile**
 - **fragment-threshold** *threshold*
 - **no fragment-threshold**
 - **lmi-type** {ansi | itu | none | rev1}
 - **mode** {dce | dte | bidir}
 - **n391dte** *intervals*
 - **no n391dte**
 - **n392dce** *threshold*
 - **no n392dce**
 - **n392dte** *threshold*
 - **no n392dte**
 - **n393dce** *count*
 - **no n393dce**
 - **n393dte** *count*
 - **no n393dte**
 - **t391dte** *keepalive*
 - **no t391dte**
 - **t392dce** *keepalive*
 - **no t392dce**
- **framing (DS3)** {c-bit | m23 | ds3-unframed}
- **idle-cycle-flag** {flags | ones}
- **no idle-cycle-flag**
- **load-balancing-algorithm** *option*
- **no load-balancing-algorithm**
- **loopback** {line | internal | remote}
- **no loopback**
- **mac** *ieee-address*
- **no mac**
- **mdl** {eic | lic | fic | unit | pfi | port | gen} *mdl-string*
- **no mdl** [eic | lic | fic | unit | pfi | port | gen]
- **mdl-transmit** {path | idle-signal | test-signal}
- **no mdl-transmit** [path | idle-signal | test-signal]
- **mode** {access | network}
- **mtu** *mtu-bytes*
- **no mtu**
- **network**
 - **accounting-policy** *policy-id*
 - **no accounting-policy**
 - [no] **collect-stats**
 - **queue-policy** *name*
 - **no queue-policy**
- **ppp**
 - **keepalive** *time-period* [**dropcount** *drop-count*]
 - **no keepalive**

- [no] **report-alarm** [ais] [los] [oof] [rai] [looped]
- [no] **scramble**
- [no] **shutdown**
- **subrate** {digital-link | larscom} *rate-step*
- no **subrate**

2.20.1.18 E1 Commands

- ```

config
 — [no] port {port-id}
 — tdm
 — [no] e1 [e1-id]
 — bert {2e3 | 2e9 | 2e11 | 2e15 | 2e20 | 2e20q | 2e23 | ones | zeros |
 alternating} duration duration
 — no bert
 — bit-error-insertion rate
 — no bit-error-insertion
 — [no] channel-group channel-group-id
 — atm
 — cell-format {uni | nni}
 — min-vp-vpi value
 — cisco-hdlc
 — down-count down-count
 — no down-count
 — keepalive time-interval
 — no keepalive
 — up-count up-count
 — no up-count
 — crc {16 | 32}
 — description long-description-string
 — no description
 — egress-scheduler-override [create]
 — no egress-scheduler-override
 — level priority-level rate pir-rate [cir cir-rate]
 — level priority-level percent-rate pir-percent [percent-cir
 cir-percent]
 — no level priority-level
 — max-rate rate
 — max-rate percent percent-rate
 — no max-rate
 — egress-scheduler-policy port-scheduler-policy-name
 — no egress-scheduler-policy
 — encap-type {atm | bcp-null | bcp-dot1q | ipcp | ppp-auto |
 frame-relay | wan-mirror | cisco-hdlc | cem}
 — frame-relay
 — [no] frf-12
 — egress
 — qos-profile profile-id
 — no qos-profile
 — fragment-threshold threshold
 — no fragment-threshold

```

---

```

— identifier frf16-identifier
— no identifier
— lmi-type {ansi | itu | none | rev1}
— mode {dce | dte | bidir}
— n391dte intervals
— no n391dte
— n392dce threshold
— no n392dce
— n392dte threshold
— no n392dte
— n393dce count
— no n393dce
— n393dte count
— no n393dte
— t391dte keepalive
— no t391dte
— t392dce keepalive
— no t392dce
— idle-cycle-flag {flags | ones}
— no idle-cycle-flag
— idle-payload-fill {all-ones}
— idle-payload-fill pattern pattern
— no idle-payload-fill
— idle-signal-fill {all-ones}
— idle-signal-fill pattern pattern
— no idle-signal-fill
— load-balancing-algorithm option
— no load-balancing-algorithm
— mac ieee-address
— no mac
— [no] mode {access | network}
— mtu mtu-bytes
— no mtu
— network
— accounting-policy policy-id
— no accounting-policy
— [no] collect-stats
— queue-policy name
— no queue-policy
— ppp
— [no] ber-sf-link-down
— compress {afc [pfc] | [pfc] afc}
— no compress
— keepalive time-period [dropcount drop count]
— no keepalive
— [no] scramble
— [no] shutdown
— speed {56 | 64}
— timeslots timeslots
— no timeslots
— clock-source {loop-timed | node-timed | adaptive | differential}
— framing (E-1) {no-crc-g704 | g704 | e1-unframed}
— [no] invert-data
— loopback {line | internal}

```

- **no loopback**
- **national-bits** *sa4 sa5 sa6 sa7 sa8*
- **no national-bits**
- **[no] report-alarm** *[ais] [los] [oof] [rai] [looped] [ber-sd] [ber-sf]*
- **[no] shutdown**
- **signal-mode** *{cas}*
- **no signal-mode** *{cas}*
- **threshold** *{ber-sd | ber-sf} rate {1 | 5 | 10 | 50 | 100}*
- **no threshold** *{ber-sd | ber-sf}*

## 2.20.1.19 E3 Commands

- ```
config
— [no] port {port-id}
— tdm
— [no] e3 [sonet-sdh-index]
— atm
— cell-format {uni | nni}
— min-vp-vpi value
— bert {2e3 | 2e9 | 2e11 | 2e15 | 2e20 | 2e20q | 2e23 | ones | zeros | alternating} duration duration
— no bert
— bit-error-insertion rate
— no bit-error-insertion
— cisco-hdlc
— down-count down-count
— no down-count
— keepalive time-interval
— no keepalive
— up-count up-count
— no up-count
— clock-source {loop-timed | node-timed}
— crc {16 | 32}
— description long-description-string
— no description
— egress-scheduler-override [create]
— no egress-scheduler-override
— level priority-level rate pir-rate [cir cir-rate]
— level priority-level percent-rate pir-percent [percent-cir cir-percent]
— no level priority-level
— max-rate rate
— max-rate percent percent-rate
— no max-rate
— egress-scheduler-policy port-scheduler-policy-name
— no egress-scheduler-policy
— encap-type {atm | bcp-null | bcp-dot1q | ipcp | ppp-auto | frame-relay | wan-mirror | cisco-hdlc | cem}
— [no] feac-loop-respond
— frame-relay
— [no] frf-12
```

```

— egress
  — qos-profile profile-id
  — no qos-profile
— fragment-threshold threshold
— no fragment-threshold
— lmi-type {ansi | itu | none | rev1}
— mode {dce | dte | bidir}
— n391dte intervals
— no n391dte
— n392dce threshold
— no n392dce
— n392dte threshold
— no n392dte
— n393dce count
— no n393dce
— n393dte count
— no n393dte
— t391dte keepalive
— no t391dte
— t392dce keepalive
— no t392dce
— framing (E-3) {g751 | g832 | e3-unframed}
— idle-cycle-flag {flags | ones}
— no idle-cycle-flag
— load-balancing-algorithm option
— no load-balancing-algorithm
— loopback {line | internal}
— no loopback
— mac ieee-address
— no mac
— mode {access | network}
— mtu mtu-bytes
— no mtu
— network
  — accounting-policy policy-id
  — no accounting-policy
  — [no] collect-stats
  — queue-policy name
  — no queue-policy
— ppp
  — keepalive time-period [dropcount drop-count]
  — no keepalive
— [no] report-alarm [ais] [los] [oof] [rai] [looped]
— [no] scramble
— [no] shutdown

```

2.20.1.20 LAG Commands

```

config
— [no] lag lag-id
  — access

```

- **adapt-qos** {link | port-fair | distribute [include-egr-hash-cfg]}
- **bandwidth** *bandwidth*
- **no bandwidth**
- **booking-factor** *factor*
- **no booking-factor**
- [no] **per-fp-egr-queuing**
- [no] **per-fp-ing-queuing**
- [no] **per-fp-sap-instance**
- **bfd**
 - [no] **disable-soft-reset-extension**
 - **family** {ipv4 | ipv6}
 - [no] **bfd-on-distributing-only**
 - **local-ip-address** *ip-address*
 - **no local-ip-address**
 - **max-admin-down-time** [[down-interval] | infinite]
 - **no max-admin-down-time**
 - **max-setup-time** [[up-interval] | infinite]
 - **no max-setup-time**
 - **multiplier** [*multiplier*]
 - **no multiplier**
 - **receive-interval** *interval*
 - **no receive-interval**
 - **remote-ip-address** *ip-address*
 - **no remote-ip-address**
 - [no] **shutdown**
 - **transmit-interval** *interval*
 - **no transmit-interval**
- **description** *long-description-string*
- **no description**
- [no] **dynamic-cost**
- **encap-type** {dot1q | null | qinq}
- **no encap-type**
- **eth-cfm**
 - [no] **mep** *mep-id* **domain** *md-index* **association** *ma-index* [vlan *vlan-id*]
 - [no] **ais-enable**
 - **client-meg-level** [*level* [*level*]]
 - **no client-meg-level**
 - [no] **interface-support-enable**
 - **interval** {1 | 60}
 - **no interval**
 - **low-priority-defect** {allDef | macRemErrXcon}
 - **priority** *priority-value*
 - **no priority**
 - **alarm-notification**
 - **fng-alarm-time** *time*
 - **fng-reset-time** *time*
 - [no] **ccm-enable**
 - **ccm-ltm-priority** *priority*
 - **no ccm-ltm-priority**
 - **ccm-padding-size** *ccm-padding*
 - **no ccm-padding-size**
 - **ccm-tlv-ignore** [port-status] [interface-status]
 - **no ccm-tlv-ignore**
 - [no] **collect-lmm-stats**

```

— [no] csf-enable
    — multiplier multiplier-value
    — no multiplier
— description description-string
— no description
— [no] eth-test-enable
    — bit-error-threshold bit-errors
    — test-pattern {all-zeros | all-ones} [crc-enable]
    — no test-pattern
— [no] facility-fault
— grace
    — eth-ed
        — max-rx-defect-window seconds
        — no max-rx-defect-window
        — priority priority
        — no priority
        — [no] rx-eth-ed
        — [no] tx-eth-ed
    — eth-vsm-grace
        — [no] rx-eth-vsm-grace
        — [no] tx-eth-vsm-grace
— low-priority-defect {allDef | macRemErrXcon | remErrXcon |
    errXcon | xcon | noXcon}
— mac-address mac-address
— no mac-address
— one-way-delay-threshold seconds
— [no] shutdown
— hold-time down hold-down-time
— no hold-time
— lacp [mode] [administrative-key admin-key] [system-id system-id] [system-priority
    priority]
— lacp-mux-control {coupled | independent}
— no lacp-mux-control
— lacp-xmit-interval {slow | fast}
— no lacp-xmit-interval
— [no] lacp-xmit-stdby
— link-map-profile lag-link-map-profile-id [create]
— no link-map-profile lag-link-map-profile-id
    — description description-string
    — no description
    — failure-mode [discard | per-link-hash]
    — no failure-mode
    — link port-id {primary | secondary}
    — no link port-id
— mac ieee-address
— no mac
— mode {access | network | hybrid}
— no mode
— [no] per-link-hash
— per-link-hash weighted [auto-rebalance]
— port port-id [port-id] [priority priority] [sub-group sub-group-id]
— no port port-id [port-id]
— port-threshold value [action {dynamic-cost | static-cost | down}] [cost static-cost]
— no port-threshold

```


- **port-type** *lag-port-type*
- **no port-type**
- **port-weight-speed** *gbps*
- **no port-weight-speed**
- **selection-criteria** [*best-port* | *highest-count* | *highest-weight*] [*slave-to-partner*]
[*subgroup-hold-time hold-time*]
- **no selection-criteria**
- [**no**] **shutdown**
- **standby-signaling** {*lACP* | *power-off*}
- **no standby-signaling**
- **weight-threshold** *value* [*action* {*dynamic-cost* | *static-cost* | *down*}] [*cost static-cost*]
- **no weight-threshold**

2.20.1.21 MACsec Commands

- ```

config
— macsec
 — connectivity-association ca-name [create]
 — no connectivity-association ca-name
 — cipher-suite {cipher-suite}
 — no cipher-suite
 — clear-tag-mode clear-tag-mode
 — no clear-tag-mode
 — description description
 — no description
 — encryption-offset encryption-offset
 — no encryption-offset
 — [no] macsec-encrypt
 — [no] replay-protection
 — replay-window-size number-of-packets
 — no replay-window-size
 — [no] shutdown
 — [no] static-cak
 — active-psk active-pre-shared-key
 — no active-psk
 — mka-hello-interval interval
 — no mka-hello-interval
 — mka-key-server-priority priority
 — no mka-key-server-priority
 — pre-shared-key pre-shared-key-index [encryption-type encryption-type] [create]
 — no pre-shared-key pre-shared-key-index
 — cak hex-string [hash | hash2]
 — no cak
 — ckn hex-string
 — no ckn
 — mac-policy mac-policy-id [create]
 — no mac-policy mac-policy-id
 — dest-mac-address mac-address [create]
 — no dest-mac-address mac-address

```

## 2.20.1.22 Ethernet Tunnel Commands

```

config
— [no] eth-tunnel tunnel-index
 — ccm-hold-time {down down-timeout | up up-timeout}
 — no ccm-hold-time
 — description long-description-string
 — no description
 — ethernet
 — encap-type {dot1q | qinq}
 — no encap-type
 — [no] mac ieee-address
 — lag-emulation
 — access
 — adapt-qos {distribute | link | port-fair}
 — no adapt-qos
 — [no] per-fp-ing-queuing
 — path-threshold num-paths
 — no path-threshold
 — [no] path path-index
 — control-tag qtag [qtag]
 — no control-tag
 — description description-string
 — no description
 — eth-cfm
 — [no] mep mep-id domain md-index association ma-index
 — alarm-notification
 — fng-alarm-time time
 — fng-reset-time time
 — [no] ccm-enable
 — ccm-ltm-priority priority
 — no ccm-ltm-priority
 — ccm-padding-size ccm-padding
 — no ccm-padding-size
 — [no] control-mep
 — description description-string
 — no description
 — [no] eth-test-enable
 — bit-error-threshold bit-errors
 — test-pattern {all zeros | all-ones} [crc-enable]
 — no test-pattern
 — grace
 — eth-ed
 — max-rx-defect-window seconds
 — no max-rx-defect-window
 — priority priority
 — no priority
 — [no] rx-eth-ed
 — [no] tx-eth-ed
 — eth-vsm-grace
 — [no] rx-eth-vsm-grace
 — [no] tx-eth-vsm-grace

```

- **low-priority-defect** {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
- **mac-address** *mac-address*
- **no mac-address**
- **one-way-delay-threshold** *seconds*
- [no] **shutdown**
- **member** *port-id*
- **no member**
- **precedence** {primary | secondary}
- **no precedence**
- [no] **shutdown**
- **protection-type** {g8031-1to1 | loadsharing}
- **revert-time** *time*
- **no revert-time**
- [no] **shutdown**

## 2.20.1.23 Multi-Chassis Redundancy Commands

- ```
config
  — redundancy
    — bgp-multi-homing
      — boot-timer seconds
      — no boot-timer
      — site-activation-timer seconds
      — no site-activation-timer
      — site-min-down-timer seconds
      — no site-min-down-timer
    — multi-chassis
      — peer ip-address [create]
      — no peer ip-address
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — description description-string
        — no description
        — [no] mc-endpoint
          — [no] bfd-enable
          — boot-timer interval
          — no boot-timer
          — hold-on-neighbor-failure multiplier
          — no hold-on-neighbor-failure
          — keep-alive-interval interval
          — no keep-alive-interval
          — [no] passive-mode
          — [no] shutdown
          — system-priority value
          — no system-priority
        — [no] mc-lag
          — hold-on-neighbor-failure multiplier
          — no hold-on-neighbor-failure
          — keep-alive-interval interval
          — no keep-alive-interval
```

-
- **lag** *lag-id lacp-key admin-key system-id system-id [remote-lag remote-lag-id] system-priority system-priority source-bmac-lsb use-lacp-key*
 - **lag** *lag-id lacp-key admin-key system-id system-id [remote-lag remote-lag-id] system-priority system-priority source-bmac-lsb MAC-Lsb*
 - **lag** *lag-id lacp-key admin-key system-id system-id [remote-lag remote-lag-id] system-priority system-priority*
 - **lag** *lag-id [remote-lag remote-lag-id]*
 - **no lag** *lag-id*
 - **[no] shutdown**
 - **mc-ring**
 - **ring** *sync-tag [create]*
 - **no ring** *sync-tag*
 - **in-band-control-path**
 - **dst-ip** *ip-address*
 - **no dst-ip**
 - **interface** *ip-int-name*
 - **no interface**
 - **service-id** *service-id*
 - **no service-id**
 - **service-name** *service-name*
 - **no service-name**
 - **[no] path-b**
 - **[no] range** *vlan-range*
 - **[no] path-excl**
 - **[no] range** *vlan-range*
 - **ring-node** *ring-node-name [create]*
 - **no ring-node** *ring-node-name*
 - **connectivity-verify**
 - **dst-ip** *ip-address*
 - **no dst-ip**
 - **interval** *interval*
 - **no interval**
 - **service-id** *service-id*
 - **no service-id**
 - **service-name** *service-name*
 - **no service-name**
 - **[no] shutdown**
 - **src-ip** *ip-address*
 - **no src-ip**
 - **src-mac** *ieee-address*
 - **no src-mac**
 - **vlan** *vlan-encap*
 - **no vlan**
 - **[no] shutdown**
 - **peer-name** *peer-name*
 - **no peer-name**
 - **[no] shutdown**
 - **source-address** *ip-address*
 - **no source-address**
 - **[no] sync**
 - **[no] igmp**
 - **[no] igmp-snooping**

- [no] **mc-ring**
- [no] **mld**
- [no] **mld-snooping**
- **pim-snooping** [saps] [spoke-sdps]
- **no pim-snooping**
- **port** *port-id* [**sync-tag** *sync-tag*] [create]
- **no port** [*port-id* | *lag-id*]
 - **range** *encap-range* **sync-tag** *sync-tag*
 - **no range** *encap-range*
- [no] **shutdown**
- [no] **srrp**
- [no] **sub-mgmt**

2.20.2 Configuration Command Descriptions

- [Generic Commands](#)
- [Card Commands](#)
- [MDA \(XMA\) Commands](#)
- [MDA/Port QoS Commands](#)
- [HSQ Port Commands](#)
- [Power Commands](#)
- [Virtual Scheduler Commands](#)
- [Forwarding Plane Configuration Commands](#)
- [General Port Commands](#)
- [Port XC Commands](#)
- [APS Commands](#)
- [Ethernet Port Commands](#)
- [802.1x Port Commands](#)
- [LLDP Port Commands](#)
- [Network Port Commands](#)
- [Interface Group Handler Commands](#)
- [Multilink-Bundle Port Commands](#)
- [SONET/SDH Port Commands](#)
- [SONET/SDH Path Commands](#)
- [ATM Interface Commands](#)
- [Frame Relay Commands](#)
- [TDM Commands](#)
- [LAG Commands](#)

- [Eth Tunnel Commands](#)
- [ETH-CFM Configuration Commands](#)
- [Multi-Chassis Redundancy Commands](#)
- [Forwarding Plane Tools Commands](#)

2.20.2.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>card>fp>ingress>access>queue-group config>card>fp>ingress>network>queue-group config>eth-tunnel>path>eth-cfm>mep config>fwd-path-ext>fpe config>lag>link-map-profile config>port>ethernet>access>egr>qgrp config>port>ethernet>access>egr>vport config>port>ethernet>access>ing>qgrp config>port>ethernet>eth-cfm>mep config>port>ethernet>network>egr>qgrp config>port>ethernet>egress>hs-sec-shaper config>port-xc>pxc config>redundancy>multi-chassis>peer
Description	This command creates a text description for a configuration context to help identify the content in the configuration file. The no form of this command removes any description string from the context. This command is supported on TDM satellite.
Default	no description
Parameters	<i>description-string</i> — Specifies a description character string. Strings can be up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

description

Syntax **description** *long-description-string*

	no description
Context	<pre> config>port config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3 config>lag </pre>
Description	This command creates a text description for a configuration context to help identify the content in the configuration file. The no form of this command removes any description string from the context.
Default	no description
Parameters	<i>long-description-string</i> — Specifies the description character string. Strings can be up to 160 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax	[no] shutdown
Context	<pre> config>card config>card>fp>ingress>mcast-path-management config>card>mcm config>card>mda config>interface-group-handler config>lag config>lag>eth-cfm>mep config>port config>port>ethernet config>port>ethernet>dwl config>port>ethernet>efm-cfm>mep config>port>ethernet>efm-oam config>port>ethernet>ssm config>port>ethernet>symbol-monitor config>port>otu config>port>sonet-sdh>path config>port>tdm>ds1 config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1 config>port>tdm>e1>channel-group config>port>tdm>e3 config>port>ml-bundle>ima>test-pattern-procedure </pre>

```

config>port>sonet-sdh>path>atm>ilmi
config>port>sonet-sdh>path>atm>ilmi
config>port>ethernet>dampening
config>port>ethernet>eth-cfm>mep
config>port>ethernet>efm-oam
config>port-xc>pxd
config>redundancy>multi-chassis>peer
config>redundancy>mc>peer>mcr
config>redundancy>mc>peer>mc-lag
config>redundancy>mc>peer>mcr>ring
config>redundancy>mc>peer>mcr>node>cv
config>redundancy>multi-chassis>peer>sync
config>redundancy>multi-chassis>peer>mc-endpoint

```

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within.

The **no** form of this command administratively enables an entity.

This command is supported on TDM satellite.

Special Cases

- card** — Specifies the default state for a card is **no shutdown**.
- interface group handler (IGH)** — Specifies the default state for an IGH is **shutdown**.
- mda** — Specifies the default state for a mda is **no shutdown**.
- lag** — Specifies the default state for a Link Aggregation Group (LAG) is **shutdown**.
- port** — Specifies the default state for a port is **shutdown**.
- path** — Specifies the default state for a SONET/SDH path is **shutdown**.

2.20.2.2 Card Commands

card

Syntax `[no] card slot-number`

Context config

Description This mandatory command enables access to the chassis card Input/Output, Control Forwarding Module (IOM/CFM), slot, MCM, MDA, XCM and XMA CLI contexts.

The **no** form of this command removes the card from the configuration. All associated ports, services, and MDAs must be shutdown.

Default	no card
Parameters	<i>slot-number</i> — Specifies the slot number of the card in the chassis. The maximum slot number is platform dependent. Refer to the hardware installation guides.
Values	1 to 10

card-type

Syntax	card-type <i>card-type</i> [level <i>card-level</i>] no card-type
Context	config>card
Description	This mandatory command adds an IOM/XCM to the device configuration for the slot. The card type can be preprovisioned, meaning that the card does not need to be installed in the chassis.

A card must be provisioned before an MDA, MCM, connector, or port can be configured.

A card can only be provisioned in a slot that is vacant, meaning no other card can be provisioned (configured) for that particular slot. To reconfigure a slot position, use the **no** form of this command to remove the current information.

A card can only be provisioned in a slot if the card type is allowed in the slot. An error message is generated if an attempt is made to provision a card type that is not allowed.

If a card is inserted that does not match the configured card type for the slot, then a log event and facility alarm is raised. The alarm is cleared when the correct card type is installed or the configuration is modified.

A log event and facility alarm are is raised if an administratively enabled card is removed from the chassis. The alarm is cleared when the correct card type is installed or the configuration is modified. A log event is issued when a card is removed that is administratively disabled.

Because IMMs do not have the capability to install separate MDAs, the configuration of the MDA is automatic. This configuration only includes the default parameters such as default buffer policies. Commands to manage the MDA such as **shutdown**, named buffer pool, and so on, remain in the MDA configuration context.

Some card hardware can support two different firmware loads. One load includes the base Ethernet functionality, including 10G WAN mode, but does not include 1588 port-based timestamping. The second load includes the base Ethernet functionality and 1588 port-based timestamping, but does not include 10G WAN mode. These are identified as two card types that are the same, except for a “-ptp” suffix to indicate the second loadset; for example, *imm40-10gb-sfp* and *imm40-10gb-sfp-ptp*. A hard reset of the card occurs when switching between the two provisioned types.

An appropriate alarm is raised if a partial or complete card failure is detected. The alarm is cleared when the error condition ceases.

New generations of cards include variants controlled by hardware and software licensing. For these cards, the license level must be provisioned in addition to the card type. A card can not become operational unless the provisioned license level matches the license level of the card installed into the slot. The set of license levels varies by card type.

The provisioned level controls aspects related to connector provisioning and the consumption of hardware egress queues and egress policers. Changes to the provisioned license level may be blocked if configuration exists that would not be permitted with the new target license level.

If the license level is not specified, the level is set to the highest license level for that card.

The **no** form of this command removes the card from the configuration.

Default no card-type

Parameters *card-type* — Specifies the type of card to be configured and installed in that slot. Values for this attribute vary by platform and release. The release notes include a listing of all supported card-types and their CLI strings. In addition, the command can be queried to check which card-types are relevant for the active platform type. Some examples include iom4-e-b and imm-2pac-fp3.

card-level — Specifies the license level of the card, up to 32 characters. Possible values vary by card type.

fail-on-error

Syntax [**no**] **fail-on-error**

Context config>card

Description This command controls the behavior of the card when any one of a specific set of card level errors is encountered in the system. When the **fail-on-error** command is enabled, and any one (or more) of the specific errors is detected, then the Operational State of the card is set to Failed. This Failed state will persist until the clear card command is issued (reset) or the card is removed and re-inserted (re-seat). If the condition persists after re-seating the card, then Nokia support should be contacted for further investigation.

Enabling **fail-on-error** is only recommended when the network is designed to be able to route traffic around a failed card (redundant cards, nodes or other paths exist).

The list of specific errors includes:

- CHASSIS event ID# 2063 – tmnxEqCardPChipMemoryEvent
- CHASSIS event ID# 2076 – tmnxEqCardPChipCamEvent
- CHASSIS event ID# 2059 – tmnxEqCardPChipError (for ingress Ethernet only)
- CHASSIS event ID# 2098 tmnxEqCardQChipBufMemoryEvent
- CHASSIS event ID# 2099 tmnxEqCardQChipStatsMemoryEvent
- CHASSIS event ID# 2101 tmnxEqCardQChipIntMemoryEvent

- CHASSIS event ID# 2103 tmnxEqCardChipIfCellEvent

On platforms without independent IOM/IMM and CPM cards, such as the 7750 SR c4/c12 or 7450 ESS-1, the node will be rebooted if fail-on-error is enabled and one of the card level errors is encountered.

The tmnxEqCardPChipError is only considered as a trigger for card fail-on-error for ingress FCS errors (not egress FCS errors), and only for Ethernet MDAs or IMMs.

Note that upon the detection of the event/error in the system, the reporting of the event (logs) and the **fail-on-error** behavior of the card are independent. Log event control configuration will determine whether the events are reported in logs (or SNMP traps, etc) and the **fail-on-error** configuration will determine the behavior of the card. This implies that the card can be configured to **fail-on-error** even if the events are suppressed (some may be suppressed in the system by default). In order to facilitate post-failure analysis, Nokia recommends that you enable the reporting of the specific events/errors (configure log event-control) when **fail-on-error** is enabled.

Default no fail-on-error

named-pool-mode

Syntax [no] named-pool-mode [now]

Context config>card

Description This command places an IOM in the named pool mode. When in named pool mode, the system will change the way default pools are created and allow for the creation of MDA and port level named buffer pools. When not enabled, the system will create default ingress and egress pools per port. When enabled, the system will not create per port pools, instead a default network and access pool is created for ingress and egress and is shared by queues on all ports.

The named pool mode may be enabled and disabled at anytime. Care should be taken when changing the pool mode for an IOM as the process of changing to or from named pool mode causes an IOM reset if MDAs are currently provisioned on the slot. If MDAs have not been provisioned at the time the named-pool-mode or no named-pool-mode command is executed, the IOM is not reset (for example, when the system is booting, the named pool mode command does not reset the IOM since the mode is set prior to provisioning the IOM's MDAs).

This command is not enabled for the ISA-AA MDA.

The **no** form of the command converts the pool mode on the IOM card to the default mode. If MDAs are currently provisioned on the IOM, the card is reset.

Default no named-pool-mode

Parameters **now** — Changes named-pool-mode without prompting.

reset-on-recoverable-error

Syntax	[no] reset-on-recoverable-error
Context	config>card
Description	This command configures the behavior of the card when a fatal memory parity error is detected on a Q-chip of the card. If reset-on-recoverable-error is enabled, the card is reset, regardless of the setting of the fail-on-error parameter. The no form of this command specifies that the recovery action is taken instead of resetting the card.
Default	no reset-on-recoverable-error

upgrade

Syntax	upgrade <i>index</i> path <i>upgrade-name</i> no upgrade <i>index</i>
Context	config>card config>card>mda
Description	This command assigns a license level upgrade to the card or XMA. There can be multiple upgrades applied to a card or XMA. The first upgrade must use index 1 and then next index 2 and so on. Also, when removing upgrades, the largest index must be removed first and then the next largest removed and so on. The path indicates the starting level and the new level that will apply due to the upgrade. For example, "cr1200g-cr1600g" can be applied to an XMA that is currently at the cr1200g level and after application of the upgrade, the operational level of the XMA shall be cr1600g. There must be an upgrade license available for the path specified. Available upgrades can be checked using the show licensing entitlements command.
Parameters	<i>index</i> — Specifies the order of the upgrade. Values 1 to 6 <i>upgrade-name</i> — Specifies the upgrade name to be applied, up to 32 characters.

2.20.2.3 MCM Commands

MCM commands are supported on the 7750 SR only.

mcm

Syntax	[no] mcm <i>mcm-slot</i>
Context	config>card
Description	This mandatory command enables access to a card's MCM CLI context to configure MCMs.
Parameters	<i>mcm-slot</i> — Specifies the MCM slot number to be configured. Even slot numbers 2 to 12 are invalid. MCM provisioning is not required to provision Compact Media Adapters (CMAs).
Values	7750 SR-c4: 1, 3 7750 SR-c12: 1, 3, 5, 7, 9, 11

mcm-type

Syntax	mcm <i>mcm-type</i> no mcm <i>mcm-type</i>
Context	config>card>mcm
Description	This mandatory command provisions a specific MCM type to the device configuration for the slot. The MCM can be preprovisioned but an MDA must be provisioned before ports can be configured. Ports can be configured once the MDA is properly provisioned. To modify an MCM slot, shut down all port associations. MCMs are required to provision MDAs. MCMs are not required to provision CMAs.
Parameters	<i>mcm-type</i> — Specifies the type of MCM to provision for that slot.
Values	mcm-v1, mcm-v2

2.20.2.4 MDA (XMA) Commands

mda

Syntax	[no] mda <i>mda-slot</i>
Context	config>card
Description	This mandatory command enables access to a card's MDA CLI context to configure MDAs.

Parameters	<i>mda-slot</i> — Specifies the MDA slot number to be configured. Slots are numbered 1 and 2. On vertically oriented slots, the top MDA slot is number 1, and the bottom MDA slot is number 2. On horizontally oriented slots, the left MDA is number 1, and the right MDA slot is number 2. For 7750 SR-c12/4 systems, MDAs may not be provisioned before MCMs are configured for the same slot. MCMs are not required for CMA provisioning.
Values	1, 2
Values	7750 SR c-12 MDA: 1, 3, 5, 7, 9, 11 7750 SR c-12 CMA: 1-12 7750 SR c-4 MDA: 1, 3 7750 SR c-4 CMA: 1-4

clock-mode

Syntax	clock-mode adaptive clock-mode differential [timestamp-freq {19440 77760 103680}]
Context	config>card>mda
Description	This command defines the clocking mode on the specified CMA/MDA. This command is only supported on CES CMAs and MDAs.
Default	clock-mode adaptive
Parameters	adaptive — Specifies that MDA is in the adaptive clock mode. This CMA/MDA can use adaptive clock recovery methods. differential — Specifies that MDA is in the differential clock mode. This CMA/MDA can use differential clock recovery methods. timestamp-freq — Sets the differential timestamp frequency to be 103.68 MHz (default), 77.76 MHz or 19.44 MHz. The frequency value is entered in kHz, thus valid values are 103680, 77760 and 19440. If this parameter is omitted, the default timestamp frequency of 103.68 MHz is used. Values 19440, 77760, 103680

fail-on-error

Syntax	[no] fail-on-error
Context	config>card>mda
Description	This command enables the fail-on-error feature. If an MDA is experiencing too many Egress XPL Errors, this feature causes the MDA to fail. This can force an APS switchover or traffic re-route . The purpose of this feature is to avoid situations where traffic is forced to use a physical link that suffers from errors but is still technically operational.

The feature uses values configured in the **config>card>mda>egress-xpl** context. When this feature is enabled on a MDA, if *window* consecutive minutes pass in which the MDA experiences more than *threshold* Egress XPL Errors per minute, then the MDA will be put in the *failed* state.

The **no** form of this command disables the feature on the MDA.

egress

Syntax	egress
Context	config>card>mda
Description	This command enables the context to configure egress MDA parameters.

egress-xpl

Syntax	egress-xpl
Context	config>card>mda
Description	This command enables the context to configure egress-xpl settings used by the fail-on-error feature.

threshold

Syntax	threshold <i>xpl-errors</i>
Context	config>card>mda>egress-xpl
Description	This command configures the Egress XPL Error Threshold value used by the fail-on-error feature.
Default	threshold 1000
Parameters	<i>xpl-errors</i> — Specifies an upper limit on the frequency of Egress XPL Errors that can occur on the MDA. When fail-on-error is enabled, if the MDA experiences more than <i>xpl-errors</i> errors per minute for the specified number of minutes from the window minutes command, the MDA will be put in the failed state. The threshold value cannot be changed while fail-on-error is enabled for this MDA.
Values	1 to 1000000

window

Syntax	window <i>minutes</i>
Context	config>card>mda>egress-xpl
Description	This command configures the Error Window value used by the fail-on-error feature.
Default	window 60
Parameters	<i>minutes</i> — Specifies the time, in minutes, that the MDA can experience frequent Egress XPL Errors. When fail-on-error is enabled, if more than <i>xpl-errors</i> Egress XPL errors per minute occur on the MDA for the specified number of consecutive minutes, the MDA will be put in the failed state. The window value cannot be changed while fail-on-error is enabled for this MDA. Values 1 to 1440

mda-type

Syntax	mda-type <i>mda-type</i> [level <i>mda-level</i>] no mda-type
Context	config>card>mda
Description	This mandatory command provisions a specific MDA type to the device configuration for the slot. The MDA can be preprovisioned but an MDA must be provisioned before ports can be configured. Ports can be configured once the MDA is properly provisioned. A maximum of two MDAs can be provisioned on an IOM/XCM. Only one MDA can be provisioned per IOM/MDA slot. To modify an MDA slot, shut down all port associations. A maximum of six MDAs or eight CMAs (or a combination) can be provisioned on a 7750 SR-c12. Only one MDA/CMA can be provisioned per MDA slot. To modify an MDA slot, shut down all port associations. CMAs do not rely on MCM configuration and are provisioned without MCMs. CMAs/XMAs are provisioned using MDA commands. A medium severity alarm is generated if an MDA/CMA is inserted that does not match the MDA/CMA type configured for the slot. This alarm is cleared when the correct MDA/CMA is inserted or the configuration is modified. A high severity alarm is raised when an administratively enabled MDA/CMA is removed from the chassis. This alarm is cleared if the either the correct MDA/CMA type is inserted or the configuration is modified. A low severity trap is issued if an MDA/CMA is removed that is administratively disabled. An MDA can only be provisioned in a slot if the MDA type is allowed in the MDA slot. An error message is generated when an MDA is provisioned in a slot where it is not allowed.

Some MDA hardware can support two different firmware loads. One load includes the base Ethernet functionality, including 10G WAN mode, but does not include 1588 port-based timestamping. The second load includes the base Ethernet functionality and 1588 port-based timestamping, but does not include 10G WAN mode. These are identified as two MDA types that are the same, except for a “-ptp” suffix to indicate the second loadset; for example, *x40-10gb-sfp* and *x40-10gb-sfp-ptp*. A hard reset of the MDA occurs when switching between the two provisioned types.

A medium severity alarm is generated if an MDA is inserted that does not match the MDA type configured for the slot. This alarm is cleared when the correct MDA is inserted or the configuration is modified.

A high severity alarm is raised when an administratively enabled MDA is removed from the chassis. This alarm is cleared if either the correct MDA type is inserted or the configuration is modified. A low severity trap is issued if an MDA is removed that is administratively disabled.

An alarm is raised if partial or complete MDA failure is detected. The alarm is cleared when the error condition ceases.

All parameters in the MDA context remain and if non-default values are required then their configuration remains as it is on all existing MDAs.

New generations of XMAs include variants controlled through hardware and software licensing. For these XMA, the license level must be provisioned in addition to the MDA type. An XMA can not become operational unless the provisioned license level matches the license level of the XMA installed into the slot. The set of license levels varies by MDA type.

The provisioned level controls aspects related to connector provisioning and the consumption of hardware egress queues and egress policers. Changes to the provisioned license level may be blocked if configuration that would not be permitted with the new target license level exists.

If the license level is not specified, the level is set to the highest license level for that XMA.

The **no** form of this command deletes the MDA from the configuration. The MDA must be administratively shut down before it can be deleted from the configuration.

Parameters *mda-type* — Specifies the type of MDA selected for the slot position. Values for this attribute vary by platform and release. The release notes include a listing of all supported mda-types and their CLI strings. In addition, the command can be queried to check which mda-types are relevant for the active platform type. Some examples include *me6-10gb-sfp+* and *x4-100g-cfp2*.

mda-level — Specifies the MDA level. Possible values vary by MDA type.

named-pool-mode

Syntax **named-pool-mode**

Context config>card>mda
config>port

Description This command enables the context to configure egress and ingress named pool parameters to store the MDA and port level named pool mode configuration commands. Currently, only the ingress and egress named-pool-policy commands are supported. Any future named pool mode configuration commands or overrides will be placed in the named-pool-mode CLI context. Within the context is an ingress and egress context.

Enter the named-pool-mode to define the ingress and egress named pool policy associations for either an MDA or port. The node may be entered regardless of the current named-pool-mode state of the IOM.

egress

Syntax egress

Context config>card>mda>named-pool-mode
config>port>named-pool-mode

Description The egress node within the named-pool-mode context is used to contain the egress named-pool-policy configuration. Enter the egress node when defining or removing the MDA or port level egress named pool policy.

ingress

Syntax ingress

Context config>card>mda>named-pool-mode
config>port>named-pool-mode

Description The ingress node within the named-pool-mode context is used to contain the ingress named-pool-policy configuration. Enter the ingress node when defining or removing the MDA or port level ingress named pool policy.

named-pool-policy

Syntax named-pool-policy *policy-name*
no named-pool-policy

Context config>card>mda>named-pool-mode>ingress
config>card>mda>named-pool-mode>egress
config>port>named-pool-mode>ingress
config>port>named-pool-mode>egress

Description	This command associates a named pool policy with an MDA or port ingress or egress context. The policy governs the way that named pools are created at the MDA or port level. The policy may be applied regardless of whether the IOM is in named pool mode; however, a named pool policy to an MDA or port to a card that is not on named pool mode will be ignored. Pools may not be created due to insufficient resources or pool name collisions. Pool name collisions are allowed. The name check is performed independently between ingress and egress. A port on ingress may have a named pool defined that is also on the egress side at the MDA level. Multiple ports on the same MDA may have the same policy or the same named pools defined. Ports on the same MDA may also have different named pool policies defined. The no named-pool-policy command removes any existing policy associated with the MDA or port.
Default	no named-pool-policy
Parameters	<i>policy-name</i> — Specifies an existing named pool policy on the system, up to 32 characters. If <i>policy-name</i> does not exist, the named-pool-policy command fails. If another named pool policy is currently associated, it continues to be associated on the MDA or port. If the policy-name does exist, the pools within the current policy (if a policy is currently associated) will be removed and the pools defined within the new policy will be created. Queues on the port or MDA will be moved to the new pools. If a pool being used by a queue existed on the previous policy, but not in the new policy, the queue is moved to the appropriate default pool and marked as 'pool-orphaned'. The policy-name may be changed at any time.

power-priority-level

Syntax	power-priority-level <i>priority</i> no power-priority-level
Context	config>card>mda
Description	This command sets the power priority value for the 7950 XRS.
Default	power-priority-level 150
Parameters	<i>priority</i> — Specifies the power priority level. An operator must assign a priority value to each XMA using a range of number from 1 to 200. The lowest number has the highest priority. The priority number range from 1 to 100 should be used for modules considered essential for system operation. Lower priority values of 101 to 200 should be used for non-essential modules.

reset-on-recoverable-error

Syntax	[no] reset-on-recoverable-error
Context	config>card>mda

Description	This command configures the behavior of the MDA when a fatal memory parity error is detected on a Q-chip of the MDA. If reset-on-recoverable-error is enabled, the MDA is reset, regardless of the setting of the fail-on-error parameter. The no form of this command specifies that the recovery action is taken instead of resetting the MDA.
Default	no reset-on-recoverable-error

sync-e

Syntax	[no] sync-e
Context	config>card>mda
Description	This command enables synchronous Ethernet on the MDA. Then any port on the MDA can be used as a source port in the sync-if-timing configuration. The no form of the command disables synchronous Ethernet on the MDA.

2.20.2.5 MDA/Port QoS Commands

access

Syntax	access
Context	config>card>mda config>port
Description	This command enables the access context to configure egress and ingress pool policy parameters. On the MDA level, access egress and ingress pools are only allocated on channelized MDAs/CMAs.

egress

Syntax	egress
Context	config>card>mda>access config>card>mda>network config>port>access config>port>network

Description This command enables the context to configure egress buffer pool parameters which define the percentage of the pool buffers that are used for CBS calculations and specify the slope policy that is configured in the **config>qos>slope-policy** context.

On the MDA level, network and access egress pools are only allocated on channelized MDAs/CMAs.

pool

Syntax [no] pool [*name*]

Context config>card>mda>access>egress
config>card>mda>access>ingress
config>card>mda>network>egress
config>port>access>egress
config>port>access>ingress
config>port>access>uplink>egress
config>port>network>egress
config>port>network>ingress

Description This command configures pool policies.

On the MDA level, access and network egress and access ingress pools are only allocated on channelized MDAs. On the MDA level, access and network egress and access ingress pools are only allocated on channelized MDAs. Network ingress pools are allocated on the MDA level for non-channelized MDAs.

Default pool default

Parameters *name* — Specifies the pool name, a string up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

amber-alarm-threshold

Syntax amber-alarm-threshold *percentage*
no amber-alarm-threshold

Context config>card>mda>access>egress>pool
config>card>mda>access>ingress>pool
config>card>mda>network>egress>pool
config>card>mda>network>ingress>pool
config>port>access>egress>pool
config>port>access>ingress>pool
config>port>network>egress>pool

Description This command configures the threshold for the amber alarm on the over-subscription allowed.

Users can selectively enable amber or red alarm thresholds. But if both are enabled (non-zero) then the red alarm threshold must be greater than the amber alarm threshold.

The **no** form of the command reverts to the default value.

Default no amber-alarm-threshold

Parameters *percentage* — Specifies the amber alarm threshold.

Values 1 to 1000

red-alarm-threshold

Syntax **red-alarm-threshold** *percentage*
no red-alarm-threshold

Context config>card>mda>access>egress>pool
config>card>mda>access>ingress>pool
config>card>mda>network>egress>pool
config>card>mda>network>ingress>pool
config>port>access>egress>pool
config>port>access>ingress>pool
config>port>network>egress>pool

Description This command configures the threshold for the red alarm on the over-subscription allowed.

Users can selectively enable amber or red alarm thresholds. But if both are enabled (non-zero) then the red alarm threshold must be greater than the amber alarm threshold.

The **no** form of the command reverts to the default value.

Default no red-alarm-threshold

Parameters *percentage* — Specifies the amber alarm threshold.

Values 1 to 1000

resv-cbs

Syntax **resv-cbs** *percent-or-default* **amber-alarm-action** **step** *percent* **max** *percent*
resv-cbs *percent-or-default*
no resv-cbs

Context config>card>mda>access>egress>pool
config>card>mda>access>ingress>pool
config>card>mda>network>egress>pool
config>card>mda>network>ingress>pool
config>port>access>egress>channel>pool
config>port>access>egress>pool

```
config>port>access>ingress>pool
config>port>ethernet>network
config>port>network>egress>pool
```

Description This command defines the percentage or specifies the sum of the pool buffers that are used as a guideline for CBS calculations for access and network ingress and egress queues. Two actions are accomplished by this command:

- A reference point is established to compare the currently assigned (provisioned) total CBS with the amount the buffer pool considers to be reserved. Based on the percentage of the pool reserved that has been provisioned, the over provisioning factor can be calculated.
- The size of the shared portion of the buffer pool is indirectly established. The shared size is important to the calculation of the instantaneous-shared-buffer-utilization and the average-shared-buffer-utilization variables used in Random Early Detection (RED) per packet slope plotting.

It is important to note that this command does not actually set aside buffers within the buffer pool for CBS reservation. The CBS value per queue only determines the point at which enqueueing packets are subject to a RED slope. Oversubscription of CBS could result in a queue operating within its CBS size and still not able to enqueue a packet due to unavailable buffers. The `resv-cbs` parameter can be changed at any time.

If the total pool size is 10 MB and the `resv-cbs` set to 5, the 'reserved size' is 500 KB.

The **no** form of this command restores the default value.

The `no resv-cbs` command will clear all the adaptive configurations. There cannot be any adaptive sizing enabled for default `resv-cbs`.

Default `resv-cbs 30`

Parameters *percent-or-default* — Specifies the pool buffer size percentage.

Values 0 to 100, or **default**

amber-alarm-action step percent — Specifies the percentage step-size for the reserved CBS size of the pool. When using the default value, the adaptive CBS sizing is disabled. To enable adaptive CBS sizing, **step percent** must be set to non-default value along with the **max** parameter. When reserved CBS is default adaptive CBS sizing cannot be enabled. The reserved CBS (Committed Burst Size) defines the amount of buffer space within the pool that is not considered shared.

Values 1 to 100

Default 0

max percent — Specifies the maximum percentage for the reserved CBS size of the pool. When using the default value, the adaptive CBS sizing is disabled. To enable adaptive CBS sizing, **max** value must be set to non-default value along with the **step percent**. When reserved CBS is default adaptive CBS sizing cannot be enabled. The reserved CBS (Committed Burst Size) defines the amount of buffer space within the pool that is not considered shared. Max reserved CBS must not be more than the reserved CBS.

Values 1 to 100

Default 0

slope-policy

Syntax	slope-policy <i>name</i> no slope-policy
Context	config>card>mda>access>egress>pool config>card>mda>access>ingress>pool config>card>mda>network>egress>pool config>card>mda>network>ingress>pool config>port>access>egress>channel>pool config>port>access>egress>pool config>port>access>ingress>pool config>port>network>egress>pool
Description	This command specifies an existing slope policy which defines high and low priority RED slope parameters and the time average factor. The policy is defined in the config>qos>slope-policy context.
Default	slope-policy default
Parameters	<i>name</i> — Specifies the policy name, a string up to 32 characters.

ingress

Syntax	ingress
Context	config>card>mda>access config>card>mda>network config>port>access
Description	This command enables the context to configure ingress buffer pool parameters which define the percentage of the pool buffers that are used for CBS calculations and specify the slope policy that is configured in the config>qos>slope-policy context. On the MDA level, access ingress pools are only allocated on channelized MDAs/CMAs.

ingress-xpl

Syntax	ingress-xpl
Context	config>card>mda
Description	This command enables the context to configure ingress MDA XPL interface error parameters.

threshold

Syntax	threshold <i>xpl-errors</i>
Context	config>card>mda>ingress-xpl
Description	This command configures the Ingress XPL Error Threshold value used by the fail-on-error feature.
Default	threshold 1000
Parameters	<i>xpl-errors</i> — Specifies an upper limit on the frequency of Ingress XPL Errors that can occur on the MDA. When fail-on-error is enabled, if the MDA experiences more than <i>xpl-errors</i> errors per minute for the specified number of minutes from the window minutes command, the MDA will be put in the failed state. The threshold value cannot be changed while fail-on-error is enabled for this MDA. Values 1 to 1000000

window

Syntax	window <i>minutes</i>
Context	config>card>mda>ingress-xpl
Description	This command configures the Error Window value used by the fail-on-error feature.
Default	window 60
Parameters	<i>minutes</i> — Specifies the time, in minutes, that the MDA can experience frequent Ingress XPL Errors. When fail-on-error is enabled, if more than <i>xpl-errors</i> Ingress XPL errors per minute occur on the MDA for the specified number of consecutive minutes, the MDA will be put in the <i>failed</i> state. The window value cannot be changed while fail-on-error is enabled for this MDA. Values 1 to 1440

network

Syntax	network
Context	config>card>mda config>port
Description	This command enables the network context to configure egress and ingress pool policy parameters. On the MDA level, network egress pools are only allocated on channelized MDAs/CMAs.

2.20.2.6 Power Commands

The following power commands are supported the 7950 XRS only.

mode

Syntax	mode {none basic advanced}
Context	config>system>power-management
Description	This command sets the power mode.
Default	mode basic
Parameters	none — Specifies that there is no management of power to modules. In this mode, no gradual shutdown of active XCMs and XMAAs is enforced. No spare capacity is reserved and any APEQ failure may result in brownouts or card failures. basic — Specifies that the node will bring up as many provisioned modules (in order of priority) as possible using the N+1 algorithm. In basic mode the system shuts down IO cards when power capacity drops below the Power Safety Level. advanced — Specifies that the operator can maintain a spare APEQ as long as possible to make it immune to the possibility of power brown-outs. In advanced mode, the system starts shutting down IO cards when the power capacity drops below the Power Safety Level + Max rated APEQ.

pcm

Syntax	[no] pcm pcm-slot [chassis chassis-id]
Context	config>system>power-management
Description	This command sets the PCM slot number.

Parameters	<i>pcm-slot</i> — Identifies the PCM slot.
Values	1 to 12
	<i>chassis-id</i> — Specifies chassis ID for the router chassis.
Values	1, 2
Default	1

pcm-type

Syntax	pcm-type <i>pcm-type</i> no pcm-type
Context	config>system>pwr-mgmt>pcm
Description	This command sets the type of PCM for the designated PCM slot. This is not a mandatory configuration; however, by configuring a PCM type of quad-pcm , this ensures the system will always monitor for the presence of PCM fan trays and will provide an indication if no PCM fan trays are detected. The no form of this command moves the PCM to an unprovisioned state.
Default	no pcm-type
Parameters	<i>pcm-type</i> — Identifies the PCM type installed in the PCM slot.
Values	dual-pcm, quad-pcm

peq

Syntax	[no] peq <i>peq-slot</i> [chassis <i>chassis-id</i>]
Context	config>system>power-management
Description	This command sets the APEQ slot number.
Parameters	<i>peq-slot</i> — Identifies the APEQ slot.
Values	1 to 12
	<i>chassis-id</i> — Specifies chassis ID for the router chassis.
Values	1, 2
Default	1

peq-type

Syntax	peq-type <i>peq-type</i>

no peq-type

Context	config>system>pwr-mgmt>peq
Description	This command sets the type of APEQ for the designated APEQ slot. The no form of this command moves the APEQ to an unprovisioned state.
Default	no peq-type
Parameters	<i>peq-type</i> — Identifies the APEQ type. Values apeq-ac-4400, apeq-ac-3000, apeq-dc-2000, apeq-dc-2200-2800, apeq-dc-4275, apeq-hvdc-3000

input-power-mode

Syntax	input-power-mode <i>amperage</i>
Context	config>system>pwr-mgmt>peq
Description	This command sets the input-power-mode of the APEQ for the designated APEQ slot.
Parameters	<i>amperage</i> — Sets the APEQ input power mode. Values 60, 80 Default 60

shutdown

Syntax	[no] shutdown
Context	config>system>pwr-mgmt>peq
Description	This command administratively enables/disables the APEQ.

power-safety-alert

Syntax	power-safety-alert <i>wattage</i>
Context	config>system>pwr-mgmt
Description	This command sets a value in watts for the Power Safety Alert. The Power Safety Alert minor alarm is generated when the system power capacity drops below the Power Safety Level (in watts) + the Power Safety Alert. This is a critical level, which when breached the system starts shutting down IO cards based on card priority.

Parameters	<i>wattage</i> — Specifies the number of Watts for the power safety alert level.
Values	0 to 102600
Default	0

power-safety-level

Syntax	power-safety-level <i>percent</i>
Context	config>system>pwr-mgmt
Description	This command sets the Power Safety Level, which is a percentage of the calculated worst case power draw value. Once a Power Safety Level is configured by the operator, both the Basic and Advanced modes use the Power Safety Level as a reference for calculating the power redundancy using N+1 algorithm during start up and recovery from power depression.
Default	power-safety-level 100
Parameters	<i>percent</i> — Specifies the Power Safety Level as a percentage of the calculated worst case power draw value.
Values	0 to 100

2.20.2.7 Virtual Scheduler Commands

virtual-scheduler-adjustment

Syntax	virtual-scheduler-adjustment
Context	config>card
Description	This command enables the context to configure the virtual scheduler processing on the card. This is only applicable to queues and to policers parented to a scheduler.

internal-scheduler-weight-mode

Syntax	internal-scheduler-weight-mode { default force-equal offered-load capped-offeredload }
	no internal-scheduler-weight-mode
Context	config>card>virt-sched-adj
Description	This command specifies the internal scheduler weight mode.
Default	default

- Parameters**
- default** — Specifies that queues are equally weighted, except at ingress for mixed-speed LAGs without **per-fp-inq-queuing** enabled and at egress for all mixed-speed LAGs, in which cases the queues are weighted based on port speed.
 - force-equal** — Specifies that the queues are always equally weighted.
 - offered-load** — Specifies that the queues are weighted based on observed offered load.
 - capped-offered-load** — Specifies that the queues are weighted based on observed offered load capped. by PIR

rate-calc-min-int

- Syntax** **rate-calc-min-int** [**fast-queue** *percent-of-default*] [**slow-queue** *percent-of-default*]
no rate-calc-min-int
- Context** config>card>virt-sched-adj
- Description** This command overrides the default minimum time that must elapse before a policer or queue's offered rate may be recalculated. A minimum time between offered rate calculations is enforced to both prevent inaccurate estimation of the offered rate and excessive input to the virtual scheduler process.

In order to smooth out rapidly fluctuating offered rates, the system averages the measured offered rate with a window of previously measured offered traffic statistics and knowledge of the time between the samples.

The window size is defined by the "rate calculation minimum interval" with offered traffic statistics being read at most four times within the window. Any previous measured offered statistics within the window are used in the averaging function. Note that if there are large numbers of samples required, for example when a large number of queues are running HQoS, then it may be that a time greater than the "rate calculation minimum interval" passes before another sample of the offered statistics can be taken for a queue. In this case, in order to calculate an offered rate, HQoS will always use two samples, the current and the previous. In this case, using a smaller **rate-calc-min-int** will have no effect on the responsiveness of HQoS to queue rate changes.

The system separates policers and queues into fast and slow categories and maintains a separate "rate calculation minimum interval" for each type. The default for each type are as follows:

Slow Queue: 1.0 seconds

Fast Queue: 0.25 seconds

The actual minimum rate calculation interval may be increased or decreased by using the fast-queue and/or slow-queue keywords (which are also applicable for policers managed by HQoS) followed by a percent value which is applied to the default interval. The default slow-queue threshold rate is 1 Mb/s. Once a policer or queue is categorized as slow, its rate must rise to 1.5 Mb/s before being categorized as a fast policer or queue. The categorization threshold may be modified by using the **slow-queue-threshold** command.

The **no** rate-calc-min-interval command restores the default fast queue and slow queue minimum rate calculation interval.

Default	no rate-calc-min-int
Parameters	<i>percent-of-default</i> — Specifies that the fast-queue percent-of-default parameter is optional and is used to modify the default minimum rate calculation time for “fast” queues. Defining 100.00 percent is equivalent to removing the override (restoring the default) on the fast queue minimum rate calculation time. Values 0.01% to 1000.00% Default 100.00% <i>percent-of-default</i> — Specifies that the slow-queue percent-of-default parameter is optional and is used to modify the default minimum rate calculation time for “slow” queues. Defining 100.00 percent is equivalent to removing the override (restoring the default) on the slow queue minimum rate calculation time. Values 0.01% to 1000.00% Default 100.00%

sched-run-min-int

Syntax	sched-run-min-int <i>percent-of-default</i> no sched-run-min-int
Context	config>card>virt-sched-adj
Description	This command overrides the default minimum time that must elapse before a virtual scheduler may redistribute bandwidth based on changes to the offered rates of member policers or queues. A minimum run interval is enforced to allow a minimum amount of “batching” queue changes before reacting to the changed rates. This minimum interval is beneficial since the periodic function of determining policer or queue offered rates is performed sequentially and the interval allows a number policer and queue rates to be determined prior to determining the distribution of bandwidth to the policers and queues. The default minimum scheduler run interval is 0.5 seconds. The sched-run-min-int command uses a percent value to modify the default interval. The no sched-run-min-int command restores the default minimum scheduler run interval for all virtual schedulers on the card.
Default	no sched-run-min-int

Parameters	<i>percent-of-default</i> — Specifies that the percent-of-default parameter is required and is used to modify the default minimum scheduler run interval for all virtual schedulers on the card. Defining 100.00 percent is equivalent to removing the override (restoring the default) for the minimum scheduler run interval.
Values	0.01% to 1000.00%
Default	100.00%

slow-queue-threshold

Syntax	slow-queue-threshold <i>kilobits-per-second</i> no slow-queue-threshold
Context	config>card>virt-sched-adj
Description	This command overrides the system default rate threshold where policers and queues are placed in the “slow” queue category. Slow rate policers and queues use a different minimum rate calculation interval time than fast rate queues. The rate is determined based on the previous calculated offered rate for the policer or queue. The default slow policer or queue rate is 1 Mb/s. The fast rate is derived by multiplying the slow rate by a factor of 1.5 resulting in a default fast rate of 1.5 Mb/s. The slow-queue-threshold command uses a “Kilobit-Per-Second” value to modify the default slow queue rate threshold and indirectly changes the fast queue rate threshold. The no version of this command restores the default slow queue and fast rate thresholds.
Default	no slow-queue-threshold
Parameters	<i>kilobits-per-second</i> — Specifies that the kilobit-per-second parameter is required and is used to modify the default slow rate threshold. Defining a value of 0 forces all policers and queues to be treated as fast rate. Defining a value of 1000 (1 Mb/s) returns the threshold to the default value and is equivalent to executing no slow-queue-threshold. The fast rate threshold is derived by multiplying the new slow rate threshold by a factor of 1.5. Values 0 to 1000000 kb/s Default 1000 kb/s

task-scheduling-int

Syntax	task-scheduling-int <i>percent-of-default</i> no task-scheduling-int
Context	config>card>virt-sched-adj

Description	This command overrides the system default time between scheduling the hierarchical virtual scheduling task. By default, the system “wakes” the virtual scheduler task every 50ms; this is equivalent to five 10ms timer ticks. The task-scheduling-int command uses a percent value parameter to modify the number of timer ticks. While the system accepts a wide range of percent values, the result is rounded to the nearest 10ms tick value. The fastest wake interval is 10ms (1 timer tick). The no scheduling-int command restores the default task scheduling interval of the card’s hierarchical virtual scheduler task.				
Parameters	<i>percent-of-default</i>: — Specifies that the percent-of-default parameter is required and is used to modify the default task scheduling interval for the hierarchical virtual scheduling task on the card. Defining 100.00 percent is equivalent to removing the override. <table><tr><td>Values</td><td>0.01% to 1000.00%</td></tr><tr><td>Default</td><td>100.00%</td></tr></table>	Values	0.01% to 1000.00%	Default	100.00%
Values	0.01% to 1000.00%				
Default	100.00%				

2.20.2.8 Forwarding Plane Configuration Commands

fp

Syntax	fp [<i>fp-number</i>]				
Context	config>card				
Description	This command enables access to the configuration of the forwarding planes on a card. The default forwarding plane is 1. When entering the fp node, if the forwarding plane number is omitted, the system will assume forwarding plane number 1. Commands can only be configured under card>fp if the hardware that the FP resides on (either a card or an XMA) is provisioned. Conversely, all commands under card>fp of the corresponding FPs are automatically removed when that hardware is unprovisioned.				
Parameters	<i>fp-number</i> — Specifies that the fp-number parameter is optional following the fp command. <table><tr><td>Values</td><td>1 to 8</td></tr><tr><td>Default</td><td>fp 1</td></tr></table>	Values	1 to 8	Default	fp 1
Values	1 to 8				
Default	fp 1				

egress

Syntax	egress

Context	config>card>fp
Description	This command enables access to the egress fp CLI context.

hs-fixed-high-thresh-delta

Syntax	hs-fixed-high-thresh-delta <i>size-in-bytes</i> no hs-fixed-high-thresh-delta
Context	config>card>fp>egress
Description	This command specifies the egress aggregate shaper high burst limit threshold delta for this HSQ IOM FP. An aggregate rate can be applied to each egress HSQ queue group, HS secondary shaper and (for subscribers configured with HS SLA expanded mode) primary shaper which manages the maximum burst limit over a specified shaping rate. Each aggregate shaper supports two thresholds which are used in conjunction with the low-burst-max-class setting. The system utilizes the lowest value attainable for each low threshold aggregate burst limit without causing shaper underrun conditions. The high burst limit threshold is determined by adding the configured hs-fixed-high-thresh-delta value to the aggregate's low burst limit threshold value. The hs-fixed-high-thresh-delta value should be set to at least two times the maximum frame size to prevent lower threshold class forwarding from also affecting the higher threshold classes when forwarding larger packet sizes. An insufficient high threshold delta defeats the intended purpose of mapping classes to the higher threshold. The hs-fixed-high-thresh-delta value can be changed at any time. Modifying the setting causes all aggregate shapers on this FP to reconfigure the low and high burst limit thresholds to reflect the new value. The no form of the command reverts this parameter to the default.
Default	hs-fixed-high-thresh-delta 4000
Parameters	<i>size-in-bytes</i> — Specifies high threshold data in bytes.
Values	0 to 65536

hs-pool-policy

Syntax	hs-pool-policy <i>name</i> no hs-pool-policy
Context	config>card>fp>egress
Description	This command specifies the HS pool policy for this FP. An HS pool policy contains the required parameters to create and size root and mid-tier buffer pools on an HSQ IOM, and apply a slope policy to each.

A single HS pool policy is supported per port FP. This command is only applicable to the HSQ IOM (iom4-e-hs) and will fail if configured on all other card types.

The **no** form of the command removes the policy and reapplies the default policy.

Default	hs-pool-policy default
Parameters	<i>name</i> — Specifies the HS pool policy name, up to 32 characters.

wred-queue-control

Syntax	wred-queue-control
Context	config>card>fp>egress
Description	This command enables the context to configure the aggregate WRED queue parameters for all WRED queues on an egress forwarding plane.

buffer-allocation

Syntax	buffer-allocation min <i>percentage</i> max <i>percentage</i> no buffer-allocation
Context	config>card>fp>egress>wred-queue-control
Description	The buffer-allocation command defines the amount of buffers that will be set aside for WRED queue buffer pools. Note that the min <i>percentage</i> and max <i>percentage</i> parameters must be set to the same value. The forwarding plane protects against cross application buffer starvation by implementing a hierarchy of buffer pools. At the top of the hierarchy are mega-pools. Mega-pools are used to manage buffers at a system application level. Two mega-pools are currently used by the system. The first (default) mega-pool services all non-WRED type queues and when WRED queues are not enabled will contain all available forwarding plane queue buffers. When WRED queuing is enabled, the second mega-pool (the WRED mega-pool) is given buffers from the default mega-pool based on the buffer-allocation command. The mega-pools provide buffers to the second tier buffer pools. The default mega-pool services all default pools and explicitly created named pools. As the name implies, the WRED mega-pool services all the WRED buffer pools created for the WRED queues. The WRED mega-pool allows each WRED queue pool to be configured to an appropriate size while allowing the sum of the WRED queue pool sizes to oversubscribe the total amount set aside for WRED queue buffering without affecting the queues using the default or named pools. No buffers are allocated to the WRED mega-pool until the wred-queue-control shutdown command is set to no shutdown. When the shutdown command is executed, all buffers allocated to the WRED mega-pool are returned to the default mega-pool and all WRED queues are returned either to their default buffer pool or their specified named buffer pool.

The **no** form of the command immediately restores the default min and max percentage values for sizing the WRED mega-pool.

Default	buffer-allocation min 25.00 max 25.00
Parameters	min <i>percentage</i> — Specifies that the required keyword defines the minimum percentage of total egress forwarding plane queue buffers that will be applied to the WRED mega-pool. The value given for <i>percentage</i> must be less than or equal to the value given for the max <i>percentage</i>. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%). Values 0.00 to 99.99 Default 25.00 max <i>percentage</i> — Specifies that the required keyword defines the maximum percentage of total egress forwarding plane queue buffers that may be applied to the WRED mega-pool. The value given for <i>percentage</i> must be equal to or greater than the value given for the min <i>percentage</i>. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%). Values 0.01 to 99.99 Default 25.00

resv-cbs

Syntax	resv-cbs min <i>percentage</i> max <i>percentage</i> no resv-cbs
Context	config>card>fp>egress>wred-queue-control
Description	This command defines the amount of buffers within the WRED mega-pool that will be set aside for WRED queues operating within their configured CBS thresholds. Note that the min <i>percentage</i> and max <i>percentage</i> parameters must be set to the same value. The forwarding plane protects against WRED queue buffer starvation by setting aside a portion of the buffers within the WRED mega-pool. The WRED queue CBS threshold defines when a WRED queue requests buffers from reserved portion of the WRED mega-pool and when it starts requesting buffers from the shared portion of the mega-pool. With proper oversubscription provisioning, this prevents a seldom active queue from being denied a buffer from the mega-pool when the shared portion of the mega-pool is congested. The WRED mega-slope reserve CBS size is controlled in the same manner as the overall sizing of the WRED mega-pool. A min and max parameter is provided to scope the range that the reserved portion based on percentages of the WRED mega-pool current size. The no form of the command immediately restores the default min and max percentage values for sizing the WRED mega-pool CBS reserve.
Default	resv-cbs min 25.00 max 25.00

Parameters	min percentage — Specifies that the required keyword defines the minimum percentage of the WRED mega-pool buffers that will be applied to the CBS reserve. The value given for <i>percentage</i> must be less than or equal to the value given for the max percentage. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%). Values 0.00 to 99.99 Default 25.00 max percentage — Specifies that the required keyword defines the maximum percentage of the IOM3-XP WRED mega-pool buffers that may be applied to the CBS reserve. The value given for <i>percentage</i> must be greater than or equal to the value given for the min percentage. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%). Values 0.01 to 99.99 Default 25.00

slope-policy

Syntax	slope-policy <i>slope-policy-name</i> no slope-policy
Context	config>card>fp>egress>wred-queue-control
Description	This command configures WRED slopes within the WRED mega-pool. The WRED slopes in the WRED mega-pool are used when WRED queues are requesting buffers from the mega-pool while they are over their CBS threshold. Once over the CBS threshold, the WRED queue stops receiving buffers from the CBS reserve in the mega-pool and starts competing for buffers in the shared portion of the mega-pool. If the packet resulting in the buffer request is inplus-profile, the packet will be associated with the highplus-slope. In-profile packets are associated with the high slope. Out-of-profile packets are associated with the low slope. Exceed-profile packets are associated with the exceed slope. While the queue is within its CBS threshold, the slopes are ignored. Within the defined slope-policy, each slope is enabled or disabled (no shutdown or shutdown) and each slope's geometry is defined as percentages of shared portion depth. If a slope is shutdown, the related traffic uses the minimum of the queue MBS and egress WRED megapool size as a drop tail. The slope-policy also defines the time average factor (TAF) value that is used to determine how the pool's weighted average depth is calculated. The higher the factor, the slower the average depth tracks the actual pool depth. The no form of the command restores the default slope policy to the WRED mega-pool. Default slope-policy default

Parameters *slope-policy-name* — Specifies which slope policy the system should apply to the WRED mega-pool. When slope-policy is not executed, the WRED mega-pool will use the default slope policy. The defined slope policy must already exist or the command will fail. 32 characters maximum.

shutdown

Syntax **[no] shutdown**

Context config>card>fp>egress>wred-queue-control

Description This command enables or disables egress WRED queue support on the forwarding plane. By default, WRED queue support is disabled (shutdown). While disabled, the various wred-queue-control commands may be executed on the forwarding plane and SAP egress QoS policies and egress queue group templates with wred-queue enabled may be applied to egress SAPs and port, respectively. The forwarding plane will allocate WRED pools to the WRED queues and the appropriate WRED mega-pool size and CBS reserve size will be calculated, but the WRED mega-pool will be empty and all buffers will be allocated to the default mega-pool. Each WRED queue will be mapped to either its appropriate default pool or an explicitly defined named pool.

Once the **no shutdown** command is executed, the calculated WRED mega-pool buffers will be moved from the default mega-pool to the WRED mega-pool. The WRED mega-pool CBS reserve size will be applied and each egress WRED queue will be moved from its default mega-pool buffer pool to its WRED pool within the WRED mega-pool hierarchy.

The **no** form of the command enables WRED queuing on an egress forwarding plane.

Default shutdown

hi-bw-mcast-src

Syntax **hi-bw-mcast-src [alarm] [group group-id] [default-paths-only]**
no hi-bw-mcast-src

Context config>card>fp

Description This command designates the forwarding plane as a high-bandwidth IP multicast source, expecting the ingress traffic to include high-bandwidth IP multicast traffic. When configured, the system attempts to allocate a dedicated multicast switch fabric plane (MSFP) to the forwarding plane. If a group is specified, all FPs in the group will share the same MSFP. If the alarm parameter is specified and the system cannot allocate a dedicated MSFP to the new group or FP, the FPs will be brought online and generate an event (SYSTEM: 2052 - tmnxChassisHiBwMulticastAlarm). Similarly, if during normal operation there is a failure or removal of resources, an event will be generated if the system cannot maintain separation of MSFPs for the MDAs.

The **no** form of the command removes the high-bandwidth IP multicast source designation from the forwarding plane.

Default	no hi-bw-mcast-src
Parameters	alarm — Enables event generation if the MDA is required to share an MSFP with another MDA that is in a different group. MDAs within the same group sharing an MSFP will not cause this alarm. <i>group-id</i> — Specifies the logical MSFP group for the MDA. MDAs configured with the same <i>group-id</i> will be placed on the same MSFP. Values 0 to 32 (A value of 0 removes the MDA from the group.) Default By default, “none” is used, and the system will attempt to assign a unique MSFP to the MDA. default-paths-only — When this parameter is specified the system will only attempt to allocate the two default paths (one high priority and one low priority) to dedicated MSFPs.

ingress

Syntax	ingress
Context	config>card>fp
Description	This command enables access to the ingress fp CLI context.

access

Syntax	access
Context	config>card>fp>ingress
Description	This CLI node contains the access forwarding-plane parameters.

queue-group

Syntax	queue-group <i>queue-group-name</i> instance <i>instance-id</i> [create] no queue-group <i>queue-group-name</i> instance <i>instance-id</i>
Context	config>card>fp>ingress>access
Description	This command creates an instance of a named queue group template on the ingress forwarding plane of a given IOM/IMM. The <i>queue-group-name</i> and instance <i>instance-id</i> are mandatory parameters when executing the command.

The named queue group template can contain only policers. If it contains queues, then the command will fail.

The **no** form of the command deletes a specific instance of a queue group.

Parameters *queue-group-name* — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM, up to 32 characters. The queue-group-name must correspond to a valid ingress queue group template name, configured under **config>qos>queue-group-templates**.

instance-id — Specifies the instance of the named queue group to be created on the IOM/IMM ingress forwarding plane.

Values 1 to 65535

create — Keyword used to associate the queue group. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>card>fp>ingress>access>queue-group
config>card>fp>ingress>network>queue-group

Description This command configures an accounting policy that can apply to a queue-group on the forwarding plane.

An accounting policy must be configured before it can be associated to an interface. If the accounting *policy-id* does not exist, an error is returned.

Accounting policies associated with service billing can only be applied to SAPs. The accounting policy can be associated with an interface at a time.

The **no** form of this command removes the accounting policy association from the queue-group.

Default No accounting policies are specified by default. You must explicitly specify a policy. If configured, the accounting policy configured as the default is used.

Parameters *acct-policy-id* — Specifies the name of the accounting policy to use for the queue-group.

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>card>fp>ingress>access>queue-group

config>card>fp>ingress>network>queue-group

Description This command enables the collection of accounting and statistical data for the queue group on the forwarding plane. When applying accounting policies, the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated, however, the CPU does not obtain the results and write them to the billing file. If the **collect-stats** command is issued again (enabled), then the counters written to the billing file will include the traffic collected while the **no collect-stats** command was in effect.

Default no collect-stats

policer-control-override

Syntax **policer-control-override [create]**
no policer-control-override

Context config>card>fp>ingress>access>queue-group
config>card>fp>ingress>network>queue-group

Description This command configures policer control overrides.

Parameters **create** — Keyword required to create a new policer control override instance.

max-rate

Syntax **max-rate {rate | max}**
no max-rate

Context config>card>fp>ingress>acc>qgrp>policer-ctrl-over
config>card>fp>ingress>network>qgrp>policer-ctrl-over

Description This command defines the parent policer's PIR leaky bucket's decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance. Packets that are not discarded by the child policers associated with the SAP or subscriber instance are evaluated against the parent policer's PIR leaky bucket.

For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet's child policer is in the confirming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet.

If the result is “conform,” the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is “violate,” the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR and FIR bandwidth.

The **policer-control-policy root max-rate** setting may be overridden on each SAP or sub-profile where the policy is applied.

The **no** version of the command returns the policer-control-policy’s parent policer maximum rate to **max**.

Default	max-rate max
Parameters	rate — Specifies that a kilobits-per-second value is mutually exclusive with the max keyword. The kilobits-per-second value must be defined as an integer that represents the number of kilobytes that the parent policer will be decremented per second. The actual decrement is performed per packet based on the time that has elapsed since the last packet associated with the parent policer. Values 0 to 2000000000 max — The max keyword is mutually exclusive with defining a kilobits-per-second value. When max is specified, the parent policer does not enforce a maximum rate on the aggregate throughput of the child policers. This is the default setting when the policer-control-policy is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

priority-mbs-thresholds

Syntax	priority-mbs-thresholds
Context	config>card>fp>ingress>access>queue-group>policer-control-override config>card>fp>ingress>network>queue-group>policer-control-override
Description	This command contains the root arbiter parent policer’s min-thresh-separation command and each priority level’s mbs-contribution command that is used to internally derive each priority level’s shared-portion and fair-portion values. The system uses each priority level’s shared-portion and fair-portion value to calculate each priority level’s discard-unfair and discard-all MBS thresholds that enforce priority sensitive rate-based discards within the root arbiter’s parent policer.

The **priority-mbs-thresholds** CLI node always exists and does not need to be created.

min-thresh-separation

Syntax	min-thresh-separation <i>size</i> [bytes kilobytes] no min-thresh-separation
Context	config>card>fp>ingress>access>queue-group>policer-control-override>priority-mbs-thresholds config>card>fp>ingress>network>queue-group>policer-control-override>priority-mbs-thresholds
Description	This command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied.

The system uses the default or specified min-thresh-separation value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the **mbs-contribution** command's optional fixed keyword is not specified):

- When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level will be zero.
- When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level will be equal to the current value of **min-thresh-separation**.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level will be zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion will be equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion will be equal to the maximum of the following:
 - **min-thresh-separation** value
 - The priority level's **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead the following rules will be used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion will be set to zero

- If the **mbs-contribution** value is not set to zero:
 - The shared-portion will be set to the current **min-thresh-separation** value
 - The fair-portion will be set to the maximum of the following:
 - **min-thresh-separation** value
 - **mbs-contribution** value less **min-thresh-separation** value

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated except for parent policer instances that currently have a min-thresh-separation override.

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that will be handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.

One thing to note is that a priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.

The **no** form of this command returns the policy's **min-thresh-separation** value to the default value. This has no effect on instances of the parent policer where **min-thresh-separation** is overridden unless the override is removed.

Default no min-thresh-separation

Parameters *size* — Specifies that the size parameter is required when executing the **min-thresh-separation** command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is the assumed value. Setting this value has no effect on parent policer instances where the **min-thresh-separation** value has been overridden. Clearing an override on parent policer instance causes this value to be enforced.

Values 0 to 16777216

bytes | kilobytes — Specifies that the **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in kilobytes.

Values bytes or kilobytes

Default kilobytes

priority

Syntax	<code>[no] priority <i>level</i></code>
Context	<code>config>card>fp>ingress>access>queue-group>policer-control-override>priority-mbs-thresholds</code> <code>config>card>fp>ingress>network>queue-group>policer-control-override>priority-mbs-thresholds</code>
Description	The priority level command contains the mbs-contribution configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority. Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.
Parameters	<i>level</i> — Specifies the priority level.
Values	1 to 8

mbs-contribution

Syntax	mbs-contribution <i>size</i> [bytes kilobytes] no mbs-contribution
Context	config>card>fp>ingress>access>queue-group>policer-control-override>priority-mbs-thresholds>priority

```
config>card>fp>ingress>network>queue-group>policer-control-override>priority-mbs-  
thresholds>priority
```

Description

This command configures the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP or subscriber context. The system uses the parent policer's **min-thresh-separation** value, the priority level's **mbs-contribution** value and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP or subscriber sub-profile where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during congestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues once all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level will be consumed by its own packets and any packets associated with higher priorities.

The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated to the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the amount of packets forwarded by the parent policer for the child's priority level. It simply modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst

tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mb/s (max-rate 20,000).
- A priority level's fair burst size is set to 30 kbytes (mbs-contribution 30 kilobytes).
- Higher priority traffic is currently taking 12 Mb/s.
- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 kbytes, which makes each child's FIR MBS 10 kbytes.
- The children want 10 Mb/s, but only 8 Mb/s is available,
- Based on weights, the children's FIR rates are set as follows:

	FIR Rate	FIR MBS
Child 1	4 Mb/s	10 kbytes
Child 2	3 Mb/s	10 kbytes
Child 3	1 Mb/s	10 kbytes

The 12 Mb/s of the higher priority traffic and the 8 Mb/s of fair traffic equal the 20 Mb/s decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mb/s of the parent policer's decrement rate, leaving 8 Mb/s of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 kbytes above 4 Mb/s,
- The burst tolerance of child 2 is based on 10 kbytes above 3 Mb/s,
- The burst tolerance of child 3 is based on 10 kbytes above 1 Mb/s.

If all three children burst simultaneously (unlikely), they will consume 30 kbytes above 8 Mb/s. This is the same as the remaining decrement rate after the higher priority traffic.

Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level's discard-all threshold is the total burst tolerance of the parent policer. In some cases the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer will be responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level's **mbs-contribution** value may be set to be greater than the sum of child policer's mbs and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority's children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer's PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, you should consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

Using the Fixed Keyword to Create Deterministic Parent Policar Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used which causes the **mbs-contribution** value to always be included in the calculation of parent policer's discard thresholds. The defined **mbs-contribution** value may be overridden on a subscriber sla-profile or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

Default no mbs-contribution

The no mbs-contribution command returns the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer will be recalculated.

Parameters	size — Specifies that the size parameter is required when executing the mbs-contribution command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level in bytes or kilobytes which is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is assumed. Setting this value has no effect on parent policer instances where the priority level's mbs-contribution value has been overridden. Clearing an override on parent policer instance causes this value to be enforced. Values 0 to 16777216 bytes, kilobytes — Specifies that the bytes keyword is optional and is mutually exclusive with the kilobytes keyword. When specified, size is interpreted as specifying the size of min-thresh-separation in bytes. The kilobytes keyword is optional and is mutually exclusive with the bytes keyword. When specified, size is interpreted as specifying the size of min-thresh-separation in kilobytes. Default kilobytes

policer-control-policy

Syntax	policer-control-policy <i>policer-control-policy-name</i> no policer-control-policy
Context	config>card>fp>ingress>access>queue-group config>card>fp>ingress>network>queue-group
Description	This command configures an policer-control policy that can apply to a queue-group on the forwarding plane. The no form of this command removes the policer-control policy association from the queue-group.
Default	no policer-control-policy
Parameters	<i>policer-control-policy-name</i> — Specifies the name of the policer-control policy to use for the queue-group. The name can be up to 32 characters long.

policer-override

Syntax	[no] policer-override
Context	config>card>fp>ingress>access>queue-group config>card>fp>ingress>network>queue-group
Description	This command, within the SAP ingress or egress contexts, is used to create a CLI node for specific overrides to one or more policers created on the SAP through the sap-ingress or sap-egress QoS policies.

The **no** form of the command removes any existing policer overrides.

Default no policer-override

policer

Syntax **policer** *policer-id* [**create**]
no policer *policer-id*

Context config>card>fp>ingress>access>qgrp>policer-over
config>card>fp>ingress>network>qgrp>policer-over

Description This command creates, modifies or deletes a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that will be associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping fails.

The **no** form of this command deletes a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscribers associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

Parameters *policer-id* — Specifies that the *policer-id* must be specified when executing the *policer* command. If the specified ID already exists, the system enters that *policer*'s context to allow the *policer*'s parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the *policer* ID will be created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new *policer* ID to the QoS policy) and the system will enter that new *policer*'s context for possible parameter modification.

Values 1 to 32

cbs

Syntax **cbs** {*size* [**bytes** | **kilobytes**] | **default**}
no cbs

Context config>card>fp>ingress>access>qgrp>policer-over>plcr
config>card>fp>ingress>network>qgrp>policer-over>plcr

Description This command configures the *policer*'s CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the *policer*. If the *policer*'s forwarding rate is equal to or less than the *policer*'s defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The *policer*'s **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command returns the *policer* to its default CBS size.

Parameters *size* — Specifies that the *size* parameter is required when specifying **cbs** and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **bytes** and **kilobytes** keywords are mutually exclusive and are used to explicitly define whether *size* represents bytes or kilobytes.

bytes — When **bytes** is defined, the value given for *size* is interpreted as the queue's CBS value given in bytes.

kilobytes — When **kilobytes** is defined, the value is interpreted as the queue's CBS value given in kilobytes.

Values 0 to 16777216

Default **kilobyte**

default — Specifying the keyword **default** sets the CBS to its default value.

mbs

Syntax **mbs** {*size* [**bytes** | **kilobyte**] | **default**}

no mbs

Context	config>card>fp>ingress>access>qgrp>policer-over>plcr config>card>fp>ingress>network>qgrp>policer-over>plcr
Description	This command configures the policer's PIR leaky bucket's high priority violate threshold. The high-prio-only command is applied to the MBS value to derive the bucket's low priority violate threshold. For ingress, trusted in-profile packets and un-trusted high priority packets use the policer's high priority violate threshold while trusted out-of-profile and un-trusted low priority packets use the policer's low priority violate threshold. At egress, in-profile packets use the policer's high priority violate threshold and out-of-profile packets use the policer's low priority violate threshold. The PIR bucket's violate threshold represent the maximum burst tolerance allowed by the policer. If the policer's offered rate is equal to or less than the policer's defined rate, the PIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the offered rate increases beyond the metering rate, the amount of data allowed above the rate is capped by the threshold. The low priority violate threshold provides a smaller burst size for the lower priority traffic associated with the policer. Since all lower priority traffic is discarded at the lower burst tolerance size, the remaining burst tolerance defined by high-prio-only is available for the higher priority traffic. The policer's mbs size defined in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. The no form of this command returns the policer to its default MBS size.
Parameters	size — The <i>size</i> parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional bytes and kilobytes keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes. bytes — When bytes is defined, the value given for size is interpreted as the policer's MBS value given in bytes. kilobytes — When kilobytes is defined, the value is interpreted as the policer's MBS value given in kilobytes. Values 0 to 16777216 default — Specifying the keyword default sets the MBS to its default value.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> } no packet-byte-offset
Context	config>card>fp>ingress>access>qgrp>policer-over>plcr config>card>fp>ingress>network>qgrp>policer-over>plcr

Description This command modifies the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism that can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** version of this command removes per packet size modifications from the policer.

Parameters *add-bytes* — The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 1 to 31

sub-bytes — The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When **b** is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet. Note that the minimum resulting packet size used by the system is 1 byte.

Values 0 to 32

rate

Syntax **rate** {*rate* | **max**} [**cir** {**max** | *rate*}]
no rate

Context config>card>fp>ingress>access>qgrp>policer-over>plcr
config>card>fp>ingress>network>qgrp>policer-over>plcr

Description This command configures the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches its exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and will be discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 kb/s (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command restores the default metering and profiling rate to a policer.

Parameters {*rate* | **max**} — Specifying the keyword **max** or an explicit *rate* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values **max** or 1 to 2000000000

cir {**max** | *rate*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *rate* parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 kb/s. The *kilobits-per-second* value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used will be equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to max.

Values **max** or 0 to 2000000000

stat-mode

Syntax	stat-mode <i>stat-mode</i> no stat mode
Context	config>card>fp>ingress>access>qgrp>policer-over>plcr config>card>fp>ingress>network>qgrp>policer-over>plcr
Description	This command configures the forwarding plane counters that allow offered, output and discard accounting to occur for the policer. An ingress policer has multiple types of offered packets (explicit in-profile, explicit out-of-profile, high priority or low priority) and each of these offered types is interacting with the policer's metering and profiling functions resulting in colored output packets (green, yellow and red). Due to the large number of policers, it is not economical to allocate counters in the forwarding plane for all possible offered packet types and output conditions. Many policers will not be configured with a CIR profiling rate and not all policers will receive explicitly profiled offered packets. The stat-mode command allows provisioning of the number of counters each policer requires and how the offered packet types and output conditions should be mapped to the counters. While a no-stats mode is supported which prevents any packet accounting, the use of the policer's parent command requires at the policer's stat-mode to be set at least to the minimal setting so that offered stats are available for the policer's Fair Information Rate (FIR) to be calculated. Once a policer has been made a child to a parent policer, the stat-mode cannot be changed to no-stats unless the policer parenting is first removed. Each time the policer's stat-mode is changed, any previous counter values are lost and any new counters are set to zero. Each mode uses a certain number of counters per policer instance that are allocated from the forwarding plane's policer counter resources. You can view the total/allocated/free stats by using the tools dump system-resources command. If insufficient counters exist to implement a mode on any policer instance, the stat-mode change will fail and the previous mode will continue unaffected for all instances of the policer. The default stat-mode when a policer is created within the policy is minimal. The stat-mode setting defined for the policer in the QoS policy may be overridden on an sla-profile or SAP where the policy is applied. If insufficient policer counter resources exist to implement the override, the stat-mode override command will fail. The previous stat-mode setting active for the policer will continue to be used by the policer. The no form of this command attempts to return the policer's stat-mode setting to minimal. The command will fail if insufficient policer counter resources exist to implement minimal where the QoS policer is currently applied and has a forwarding class mapping.
Parameters	— See the <i>7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide</i> for details on the policer stat-mode parameters.

mcast-path-management

Syntax	mcast-path-management
Context	config>card>fp>ingress
Description	This CLI node contains the forwarding plane settings for ingress multicast path management. Enter the node to configure the bandwidth-policy and the administrative state of ingress multicast path management.

bandwidth-policy

Syntax	bandwidth-policy <i>policy-name</i> no bandwidth-policy				
Context	config>card>fp>ingress>mcast-path-management				
Description	This command explicitly associates a bandwidth policy to a forwarding plane or MDA. The bandwidth policy defines the dynamic rate table and the multicast paths bandwidth and queuing parameters. If a bandwidth policy is not explicitly associated with a forwarding plane or MDA, the default bandwidth policy is used when ingress multicast path management is enabled. The no form of the command removes an explicit bandwidth policy from a forwarding plane or MDA and restores the default bandwidth policy.				
Parameters	<i>policy-name</i> — The policy-name parameter is required and defines the bandwidth policy that should be associated with the MDA or forwarding plane for ingress multicast path management. If the policy name does not exist, the bandwidth-policy command will fail. The name can be up to 32 characters long. <table><tr><td>Values</td><td>Any existing bandwidth policy name.</td></tr><tr><td>Default</td><td>default</td></tr></table>	Values	Any existing bandwidth policy name.	Default	default
Values	Any existing bandwidth policy name.				
Default	default				

queue-group

Syntax	queue-group <i>queue-group-name</i> instance <i>instance-id</i> [create] no queue-group <i>queue-group-name</i> instance <i>instance-id</i>
Context	config>card>fp>ingress>network
Description	This command creates a queue-group instance in the network ingress context of a forwarding plane. Only a queue-group containing policers can be instantiated. If the queue-group template contains policers and queues, the queues are not instantiated. If the queue-group contains queues only, the instantiation in the data path is failed.

One or more instances of the same policer queue-group name and/or a different policer queue-group name can be created on the network ingress context of a forwarding plane.

The queue-group-name must be unique within all network ingress and access ingress queue groups in the system. The queue-group instance-id must be unique within the context of the forwarding plane.

The **no** version of this command deletes the queue-group instance from the network ingress context of the forwarding plane.

Parameters *queue-group-name* — Specifies the name of the queue group template up to 32 characters.

instance-id — Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

create — Keyword used to create the queue-group instance.

ingress-buffer-allocation

Syntax `ingress-buffer-allocation percentage`
`no ingress-buffer-allocation`

Context config>card>fp

Description	This command allows the user to configure an ingress buffer allocation percentage per forwarding plane from 20.00% to 80.00%. Ingress buffer allocation applies to user-accessible buffers (total buffers less those reserved for system use).

The ingress buffer allocation percentage determines how much of the user-accessible buffers will be available for ingress purposes. The remaining buffers will be available for egress purposes.

This command is supported on all 50G FP2-based and 100G/200G FP3-based hardware. It is not supported on other FP2 or FP3-based hardware, nor on FP4-based hardware.

The **no** form of this command returns the ingress buffer allocation to the default value.

Default ingress-buffer-allocation 50.00

Parameters *percentage* — Specifies the buffer allocation percentage.

Values 20.00 to 80.00

policy-accounting

Syntax **policy-accounting** *limit*
 no policy-accounting

Context	config>card>fp
Description	This command allows the user to configure the number of stats resources for policy accounting for the forwarding plane.
Default	no policy-accounting
Parameters	<i>limit</i> — Specifies the number of stats resources.
Values	1000 to 128000

stable-pool-sizing

Syntax	[no] stable-pool-sizing
Context	config>card>fp
Description	This command provides a stable buffer pool allocation environment for all default port buffer pools on a forwarding plane. This stable environment is provided at the expense of optimal buffer allocation between the various port buffer pools. Normally, port pools are sized according to a ports relative bandwidth with other ports and the ability of a port to use pool buffers. As an example, on a forwarding plane with two potential MDAs and only one equipped, the normal behavior is to provide all available default pool buffers to the ports on the currently equipped MDA. If a second MDA is equipped in the future, buffers are freed from the existing MDA and provided to the ports on the new MDA. Stable pool sizing alters this behavior by reserving buffers for both MDAs whether they are equipped or not thus preventing a resizing event when an MDA is equipped. In addition, existing ports on a module always receive their maximum bandwidth share of buffers independent on any sub-rate condition that may currently exist. This provides a stable amount of buffers to other ports on the module independent of link or configuration events that may occur on the port.

Stable pool sizing preserves the ability to modify the effective bandwidth used to determine a port's relative share of the available buffers through the use of the ing-percentage-of-rate and egr-percentage-of-rate commands under the port configuration. Changing the values associated with these commands will cause a reevaluation of buffer distribution and thus a possible resizing of pools on each port within the module. These commands have no effect on ports associated with other modules on the forwarding plane.

Stable pool sizing is mutually exclusive with card level named-pool-mode. Named pool mode must be disabled and not operational before stable pool sizing can be enabled. Once stable pool sizing is enabled on any forwarding plane on a card, named-pool-mode cannot be enabled for that card.

Stable pool sizing may be enabled (while named pool mode is disabled) or disabled at any time on a forwarding plane. The system will dynamically change the pool sizes according to the stable pool sizing state.

When a port connector breakout is configured, its ports will be included in the stable pool sizing calculation. Consequently, adding or removing a port connector breakout to or from the configuration will cause the buffer pool allocation to be recalculated even when stable pool sizing is enabled.

The **no** stable-pool-sizing command disables stable pool sizing on a forwarding plane. Existing buffer pools will be resized according to normal pool sizing behavior.

2.20.2.9 MACsec Commands

macsec

Syntax	macsec
Context	config
Description	This command enables the context for MACsec configuration. The MACsec MKA profile can be configured under this command.

connectivity-association

Syntax	connectivity-association <i>ca-name</i> [create] no connectivity-association <i>ca-name</i>
Context	config>macsec
Description	This command configures a connectivity association. MACsec connectivity associations are applied to a port dot1x configuration to enable MACsec on that port. The no form of this command removes the connectivity association.
Parameters	<i>ca-name</i> — The name of the connectivity association, a string up to 32 characters long. create — Mandatory while creating an entry.

cipher-suite

Syntax	cipher-suite <i>cipher-suite</i> no cipher-suite
Context	config>macsec>connectivity-association
Description	This command configures encryption of datapath PDUs. When all parties in the Connectivity Association (CA) have the SAK, they use the above algorithm in conjunction with the SAK to encrypt the datapath PDUs.

The XPN 64 bit (extended packet number) can be used for higher rate ports such as 10 GigE to minimize the window rollover and renegotiation of the SAK.

The **no** form of this command disables encryption of datapath PDUs.

Default cipher-suite gcm-aes-128

Parameters *cypher-suite* — Specifies the algorithm.

Values

- gcm-aes-128 — algorithm is used for control plain encryption
- gcm-aes-256 — algorithm is used for control plain encryption
- gcm-aes-xpn-128 — algorithm with extended packet number is used for control plain encryption
- gcm-aes-xpn-256 — algorithm with extended packet number is used for control plain encryption

clear-tag-mode

Syntax **clear-tag-mode** *clear-tag-mode*
no clear-tag-mode

Context config>macsec>connectivity-association

Description This command puts 802.1Q tags in clear "before sectag". There are two modes: single tag and dual-tag.

[Table 42](#) explains the encrypted dot1q and QinQ packet format when clear-tag-mode single-tag or dual-tag is configured.

The **no** form of this command will put all dot1q tags after sectag and encrypt the tags.

Table 42 Encrypted Dot1q and QinQ Packet Format

Unencrypted format	Clear-tag-mode	Pre-encryption (Tx)	Pre-decryption (Rx)
Single tag (dot1q)	single-tag	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTag
Single tag (dot1q)	dual-tag	DA, SA, TPID, VID, Etype	DA, SA, TPID, VID, SecTag
Double tag (q-in-q)	single-tag	DA, SA, TPID1, VID1, IPID2, VID2, Etype	DA, SA, TPID1, VID1, SecTag

Table 42 Encrypted Dot1q and QinQ Packet Format (Continued)

Unencrypted format	Clear-tag-mode	Pre-encryption (Tx)	Pre-decryption (Rx)
Double tag (QinQ)	dual-tag	DA, SA, TPID1, VID1, IPID2, VID2, Etype	DA, SA, TPID1, VID1, IPID2, VID2, SecTag

Default no clear-tag-mode

Parameters *clear-tag-mode* — Specifies the clear tag mode.

Values single-tag, dual-tag

description

Syntax **description** *description*
no description

Context config>macsec>connectivity-association

Description This command enters a description for connectivity association.

The **no** form of this command removes the connectivity association description.

Parameters *description* — The brief explanation of the connectivity association, a string up to 80 characters long.

encryption-offset

Syntax **encryption-offset** *encryption-offset*
no encryption-offset

Context config>macsec>connectivity-association

Description This command specifies the offset of the encryption in MACsec packet.

The encryption-offset is distributed by MKA (Key-server) to all parties.

It is signaled via MACsec capabilities. There are four basic settings for this. [Table 43](#) breaks down the settings.

Table 43 **MACsec Basic Settings**

Setting	Description
0	MACsec is not implemented
1	Integrity without confidentiality
2	The following are supported: • Integrity without confidentiality • Integrity and confidentiality with a confidentiality offset of 0
3 ¹	The following are supported: • Integrity without confidentiality • Integrity and confidentiality with a confidentiality offset of 0, 30, or 50

Note:

1. SR OS supports (3) Integrity without confidentiality and Integrity and confidentiality with a confidentiality offset of 0, 30, or 50.

The **no** form of this command rejects all arriving traffic whether MACsec is secured or not.

Default encryption-offset 0

Parameters *encryption-offset* — Specifies the encryption.

Values 0 — encrypt the entire payload
 30 — leave the IPv4 header in clear
 50 — leave the IPv6 header in clear

macsec-encrypt

Syntax [no] **macsec-encrypt**

Context config>macsec>connectivity-association

Description This command specifies all PDUs will be encrypted and authenticated (ICV payload).

 The **no** form of this command specifies all PDUs are transmitted with clear text, but still authenticated and have the trailing ICV.

Default macsec-encrypt

replay-protection

Syntax	[no] replay-protection
Context	config>macsec>connectivity-association
Description	Specifies the size of the replay protection window. This command must be configured to force packet discard when it has detected a packet that is not within the replay-window-size. When replay protection is enabled, the sequence of the ID number of the received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay window size, the packet is counted by the receiving port and then discarded. For example, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is counted and discarded because it falls outside the parameters of the replay window size. Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network. Replay protection should not be enabled in cases where packets are expected to arrive out of order.
Default	replay-protection

replay-window-size

Syntax	replay-window-size <i>number-of-packets</i> no replay-window-size
Context	config>macsec>connectivity-association
Description	This command specifies the size of the replay protection window. This command must be configured to enable replay protection. When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving port. For example, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

When the *number-of-packets* variable is set to 0, all packets that arrive out-of-order are dropped.

The **no** form of this command reverts to the default.

Default	replay-window-size 0
Parameters	<i>number-of-packets</i> — Specifies the window that the packets can arrive out of order.
Values	0 to 4294967294

shutdown

Syntax	[no] shutdown
Context	config>macsec>connectivity-association
Description	This command shuts down the CA profile. All ports using this profile will not transmit PDUs as this command shuts down the MACsec for this profile.
Default	shutdown

static-cak

Syntax	[no] static-cak
Context	config>macsec>connectivity-association
Description	This command allows the configuration of a Connectivity Association Key (CAK). A CAK is responsible for managing the MKA.

active-psk

Syntax	active-psk <i>active-pre-shared-key</i> no active-psk
Context	config>macsec>conn-assoc>static-cak

Description	This command specifies which pre-shared-key is the active transmitting pre-shared-key. If there are two pre-shared-keys configured, the arriving MACsec MKA can be decrypted via CAKs of both pre-shared keys; however, only the active-psk will be used for TX encryption of MKA PDUs.
Default	active-psk 1
Parameters	<i>active-pre-shared-key</i> — Specifies the value of the pre-shared-key.
Values	1 or 2

mka-hello-interval

Syntax	mka-hello-interval <i>interval</i> no mka-hello-interval
Context	config>macsec>conn-associ>static-cak
Description	This command specifies the MKA hello interval. The no form of this command disables the MKA hello interval.
Default	mka-hello-interval 2
Parameters	<i>interval</i> — Specifies the MKA hello interval, in seconds.
Values	1 to 60

mka-key-server-priority

Syntax	mka-key-server-priority <i>priority</i> no mka-key-server-priority
Context	config>macsec>conn-associ>static-cak
Description	This command specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode. The no form of this command disables the mka-key-server-priority .
Default	mka-key-server-priority 16
Parameters	<i>priority</i> — Specifies the priority of the server.
Values	0 to 255

pre-shared-key

Syntax	pre-shared-key <i>pre-shared-key-index</i> [encryption-type <i>encryption-type</i>] [create] no pre-shared-key <i>pre-shared-key-index</i>
Context	config>macsec>conn-assoc>static-cak
Description	This command specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode. This command also specifies the encryption algorithm used for encrypting the SAK. A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). The pre-shared key-the CKN and CAK-must match on both ends of a link. A pre-shared key is configured on both devices at each end of point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the successful MKA liveness negotiation. The encryption-type is used for encrypting SAK and authentication of the MKA packet. The symmetric encryption key SAK (Security Association Key) needs to be encrypted (wrapped) via the above protocols. The AES key is derived via pre-shared-key. The no form of this command removes the index.
Parameters	<i>pre-shared-key-index</i> — Specifies the index of this pre-shared-key. Values 1, 2 <i>encryption-type</i> — Specifies the type of encryption. Values aes-128-cmac, aes-256-cmac create — Mandatory while creating an entry.

cak

Syntax	cak <i>hex-string</i> [hash hash2] no cak
Context	config>macsec>conn-assoc>static-cak>pre-shared-key
Description	Specifies the connectivity association key (CAK) for a pre-shared key. Two values are derived from CAK. • KEK (Key Encryption Key), this is used to encrypt the MKA and SAK (symmetric key used for datapath PDUs) to be distributed between all members. • ICK (Integrity Check Value), this is used to authenticate the MKA and SAK PDUs to be distributed between all members. The no form of this command removes the value.

- Parameters** *hex-string* — Specifies the value of the CAK.
- Values** up to 64 hexadecimal characters, 32 hexadecimal characters for 128-bit key and 64 hexadecimal characters for 256-bit key
- hash** — Keyword, specifying the hash scheme.
- hash2** — Keyword, specifying the hash scheme.

ckn

- Syntax** **ckn** *hex-string*
no ckn
- Context** config>macsec>conn-assoc>static-cak>pre-shared-key
- Description** Specifies the connectivity association key name (CKN) for a pre-shared key.
CKN is appended to the MKA for identification of the appropriate CAK by the peer.
The **no** form of this command.
- Parameters** *hex-string* — Specifies the value of the CKN.
- Values** 32 octets char (64 hex)

mac-policy

- Syntax** **mac-policy** *mac-policy-id* [**create**]
no mac-policy *mac-policy-id*
- Context** config>macsec
- Description** This command configures MAC address policy groups.
The **no** form of this command removes the MAC address policy group configuration.
- Parameters** *mac-policy-id* — Specifies the value of the MAC address policy.
- Values** 0 to 4294967295
- create** — Mandatory to create the configuration.

dest-mac-address

- Syntax** **dest-mac-address** *mac-address* [**create**]
no dest-mac-address *mac-address*
- Context** config>macsec>mac-policy

Description	This command specifies the destination MAC address.
	The no form of this command removes the MAC address.
Parameters	<i>mac-address</i> — Specifies the value of the MAC address policy.
	Values <i>xx:xx:xx:xx:xx:xx</i> or <i>xx-xx-xx-xx-xx-xx</i>
	create — Mandatory to create the configuration.

2.20.2.10 General Port Commands

port

Syntax	[no] port { <i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>aps-id</i> <i>connector-port-id</i> }												
Context	config												
Description	This command enables access to the context to configure ports, multilink bundles, and bundle protection groups (BPGs). Before a port can be configured, the chassis slot must be provisioned with a valid card type and the MDA parameter must be provisioned with a valid MDA type.												
Default	No ports are configured. All ports must be explicitly configured and enabled.												
Parameters	<i>port-id</i> — Specifies the physical port ID in the following format: Values <i>slot/mda/port</i> [.channel]												
	<i>eth-sat-id</i> — Specifies the Ethernet satellite ID to be associated with this IP interface. This parameter applies to the 7950 XRS only. Values <table><tr><td>eth-sat-id</td><td>esat-id/slot/port</td><td></td></tr><tr><td></td><td>esat</td><td>keyword</td></tr><tr><td></td><td>id</td><td>1 to 20</td></tr></table>	eth-sat-id	esat-id/slot/port			esat	keyword		id	1 to 20			
eth-sat-id	esat-id/slot/port												
	esat	keyword											
	id	1 to 20											
	<i>pxc-id</i> — Specifies the PXC ID to be associated with this IP interface. This parameter applies to the 7950 XRS only. Values <table><tr><td>pxc-id</td><td>pxc-id.sub-port</td><td></td></tr><tr><td></td><td>pxc</td><td>keyword</td></tr><tr><td></td><td>id</td><td>1 to 64</td></tr><tr><td></td><td>sub-port</td><td>a, b</td></tr></table>	pxc-id	pxc-id.sub-port			pxc	keyword		id	1 to 64		sub-port	a, b
pxc-id	pxc-id.sub-port												
	pxc	keyword											
	id	1 to 64											
	sub-port	a, b											

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The command syntax must be configured as follows:

Values *bundle-type-slot/mda.bundle-num*

bundle-ppp-slot/mda.bundle-num	Creates a multilink PPP bundle.
bundle-ima-slot/mda.bundle-num	Creates an IMA bundle.
bundle-fr-slot/mda.bundle-num	Creates an MLFR bundle.

where: **bundle:** keyword
 slot: IOM/MDA slot numbers
 bundle-num: 1 to 336

For example:

```
router1>config# port bundle-ppp-5/1.1 (multilink PPP bundle)
router1>config# port bundle-ima-5/1.2 (IMA bundle)
```

aps-id — This option configures APS on unbundled SONET/SDH ports. All SONET-SDH port parameters, with certain exceptions, for the working and protection circuit ports must be configured in the **config>port>aps-group-id** context. The working and protection circuit ports inherit all those parameters configured. The exception parameters for the working and protect circuits can be configured in the **config>port>sonet-sdh** context. Exception list commands include:

- clock-source
- [no] loopback
- [no] report-alarm
- section-trace
- [no] threshold

When an **aps-group-id** is created all applicable parameters under the port CLI tree (including parameters under any submenus) assume **aps-group-id** defaults, or when those are not explicitly specified, default to SONET/SDH port defaults for any SONET port.

All but a few exception SONET/SDH parameters for the working channel port must be configured in the **config>port>aps>sonet-sdh** context. The protection channel inherits all the configured parameters. The exception parameters for the protection channel can be configured in the **config>port>aps>sonet-sdh** context.

Signal failure (SF) and signal degrade (SD) alarms are not enabled by default on POS interfaces. It is recommended to change the default alarm notification configuration for POS ports that belong to APS groups in order to be notified of SF/SD occurrences to be able to interpret the cause for an APS group to switch the active line.

For path alarms, modify the logical line *aps-id* in the **config>port aps-id <sonet-sdh>path report-alarm** context. For example:

```
configure port aps-1 sonet-sdh path report-alarm p-ais
```

For line alarms, separately, modify the 2 physical ports that are members of the logical aps-id port (the working and protect lines). APS reacts only to line alarms, not path alarms.

For example:

configure port **1/2/3** sonet-sdh report-alarm lb2er-sd

configure port **4/5/6** sonet-sdh report-alarm lb2er-sd

If the SD and SF threshold rates must be modified, the changes must be performed at the line level on both the working and protect APS port member.

The **no** form of this command deletes an *aps-group-id* or bundle-aps-group-id. In order for an *aps-group-id* to be deleted,

The same rules apply for physical ports, bundles deletions apply to APS ports/ bundles deletions (for example an *aps-group-id* must be shutdown, have no service configuration on it, and no path configuration on it). In addition working and protection circuits must be removed before an *aps-group-id* may be removed.

Values **port** *aps-group-id*
 aps:keyword
 where *group-id*: 1 to 64

Example: **port aps-64**

bpgrp-id — Creates a bundle protection group (BPG). The BPG consists of a working and protection bundles that provide APS protection to each other using bi-directional APS as supported on the 7750 SR family of products. All members of a working/ protection bundle must be on the same working/protection circuit respectively of the same, already provisioned APS group. The working bundle must have already been created in the **config>port** context before services can be created on a BPG.

Values **bpgrp-type-bpgrp-num**

bpgrp: keyword
type: **ppp** — Provides protection of one PPP bundle by another.
 ima — Provides protection of one IMA bundle by another IMA bundle.
bpgrp-num: 1 to 1600

connector-port-id — Specifies the physical port of a connector in the following format.

Values *slot/mda/cconnector/port*

connector

Syntax	connector
Context	config>port
Description	This command enables the context to configure connector parameters.

breakout

Syntax	breakout <i>breakout</i> no breakout
Context	config>port>connector
Description	This command defines the port breakout of the transceiver that will be used in the connector. Specifying the type triggers the creation of the ports that will be accessible under the connector. The no version of this command removes the ports under the connector.
Default	no breakout
Parameters	<i>breakout</i> — Specifies the breakout type. Values c1-40g, c4-10g, c1-100g

ddm-events

Syntax	[no] ddm-events
Context	config>port
Description	This command enables Digital Diagnostic Monitoring (DDM) events for the port. The no form of the command disables DDM events.

dwdm

Syntax	dwdm
Context	config>port
Description	This command configures the Dense Wavelength Division Multiplexing (DWDM) parameters.

amplifier

Syntax	amplifier
Context	config>port>dwdm
Description	This command enables you to tune the optical amplifier parameters.

report-alarms

Syntax	[no] report-alarms [ild] [tmp] [mth] [mtl] [los] [lop] [com]
Context	config>port>dwdm>amplifier
Description	This command allows users to enable/disable the optical amplifier alarms for the port.
Default	All alarms are enabled
Parameters	ild — Reports amplifier pump over-current faults. tmp — Reports pump temperature faults. mth — Reports module case temperature high faults. mtl — Reports module case temperature low faults. los — Reports loss of signal faults. lop — Reports loss of optical power faults. com — Reports module communication failure faults.

coherent

Syntax	coherent
Context	config>port>dwdm
Description	This command configures the coherent optical module parameters.

channel

Syntax	channel <i>channel</i>
Context	config>port>dwdm config>port>dwdm>coherent config>port>dwdm>tdcm
Description	This command configures the Dense Wavelength Division Multiplexing (DWDM) ITU channel at which a tunable MDA optical interface will be configured to operate. It is expressed in a form that is derived from the laser's operational frequency. For example 193.40 THz corresponds to DWDM ITU channel 34 in the 100 GHz grid and 193.45 THz corresponds to DWDM ITU channel 345 in the 50 GHz grid. Provisioning rules: The provisioned MDA type must have DWDM tunable optics (for example, p1-100g-tun). • The 'dwdm channel' must set to a non-zero value before the port is set to 'no shutdown' • The port must be 'shutdown' before changing the dwdm channel. • The port must be a physical port to set the dwdm channel

Parameters *channel* — Specifies the channel.

Values 0, 17 to 61, 175 to 605

Where: 17 to 61 is used for 100 GHz channels
175, 185 to 605 is used for 50 GHz channels
0 is only valid on disabled (shutdown) ports

Values The DWDM channel number range is listed in [Table 44](#).

Table 44 DWDM Channel Numbers

C-Band					
100 GHz Grid			50GHz Grid		
nm	THz	ITU Channel	nm	THz	ITU Channel
1528.77	196.10	61	1529.16	196.05	605
1529.55	196.00	60	1529.94	195.95	595
1530.33	195.90	59	1530.72	195.85	585
1531.12	195.80	58	1531.51	195.75	575
1531.90	195.70	57	1532.29	195.65	565
1532.68	195.60	56	1533.07	195.55	555
1533.47	195.50	55	1533.86	195.45	545
1534.25	195.40	54	1534.64	195.35	535
1535.04	195.30	53	1535.43	195.25	525
1535.82	195.20	52	1536.22	195.15	515
1536.61	195.10	51	1537.00	195.05	505
1537.40	195.00	50	1537.79	194.95	495
1538.19	194.90	49	1538.58	194.85	485
1538.98	194.80	48	1539.37	194.75	475
1539.77	194.70	47	1540.16	194.65	465
1540.56	194.60	46	1540.95	194.55	455
1541.35	194.50	45	1541.75	194.45	445
1542.14	194.40	44	1542.54	194.35	435

Table 44 DWDM Channel Numbers (Continued)

C-Band					
100 GHz Grid			50GHz Grid		
nm	THz	ITU Channel	nm	THz	ITU Channel
1542.94	194.30	43	1543.33	194.25	425
1543.73	194.20	42	1544.13	194.15	415
1544.53	194.10	41	1544.92	194.05	405
1545.32	194.00	40	1545.72	193.95	395
1546.12	193.90	39	1546.52	193.85	385
1546.92	193.80	38	1547.32	193.75	375
1547.72	193.70	37	1548.11	193.65	365
1548.51	193.60	36	1548.91	193.55	355
1549.32	193.50	35	1549.72	193.45	345
1550.12	193.40	34	1550.52	193.35	335
1550.92	193.30	33	1551.32	193.25	325
1551.72	193.20	32	1552.12	193.15	315
1552.52	193.10	31	1552.93	193.05	305
1553.33	193.00	30	1553.73	192.95	295
1554.13	192.90	29	1554.54	192.85	285
1554.94	192.80	28	1555.34	192.75	275
1555.75	192.70	27	1556.15	192.65	265
1556.55	192.60	26	1556.96	192.55	255
1557.36	192.50	25	1557.77	192.45	245
1558.17	192.40	24	1558.58	192.35	235
1558.98	192.30	23	1559.39	192.25	225
1559.79	192.20	22	1560.20	192.15	215
1560.61	192.10	21	1561.01	192.05	205
1561.42	192.00	20	1561.83	191.95	195
1562.23	191.90	19	1562.64	191.85	185

Table 44 DWDM Channel Numbers (Continued)

C-Band					
100 GHz Grid			50GHz Grid		
nm	THz	ITU Channel	nm	THz	ITU Channel
1563.05	191.80	18	1563.45	191.75	175
1563.86	191.70	17	—	—	—

compatibility

Syntax	compatibility <i>mode</i>
Context	config>port>dwdm>coherent
Description	This command configures the optical mode and rate of operation.
Parameters	<i>mode</i> — Specifies the optical mode. Values long-haul - The port will operate in the native long haul mode. metro - The port will operate in the native metro regional mode. access - The port will operate in the native access mode (80km reach). interop - The port will operate in the third party interop mode. interop2-The port will operate in the third party interop mode with alternate differential encoding. interop3-The port will operate in the CFP2-DCO Rev A0 Staircase FEC interop mode. Default long-haul

cpr-window-size

Syntax	cpr-window-size <i>window-size</i>
Context	config>port>dwdm>coherent
Description	This command configure the window size used for carrier phase recovery.
Default	32

Parameters *window-size* — Indicates the number of symbols used for carrier phase recovery algorithm of the receiver. When this parameter is changed, the link will bounce because the receiver needs to be reconfigured.

Values 2, 4, 8, 16, 32, 64

dispersion

Syntax **dispersion** *dispersion*

Context config>port>dwdm>coherent

Description This command configures the residual chromatic dispersion to be compensated when the coherent receiver is operating in manual dispersion control mode.

Default 0

Parameters *dispersion* — Specifies the dispersion compensation.

Values -50000 to 50000

mode

Syntax **mode** {**automatic** | **manual**}

Context config>port>dwdm>coherent

Description This command configures the mode used to compensate for chromatic dispersion.

Parameters **automatic** — Sets to automatic mode.

manual — Sets to manual mode.

report-alarms

Syntax [**no**] **report-alarms** [**modflt**] [**mod**] [**netrx**] [**nettx**] [**hosttx**]

Context config>port>dwdm>coherent

Description This command configures the alarms that will be reported for the coherent module.

Default modflt mod netrx nettx hosttx

Parameters **modflt** — Reports module fault alarm.

mod — Reports module alarm.

netrx — Reports network (optical side) receive alarm.

nettx — Reports network (optical side) transmit alarm.

hosttx — Reports host (electrical side) transmit alarm.

rx-los-reaction

Syntax	rx-los-reaction { squelch } no rx-los-reaction
Context	config>port>dwdm>coherent
Description	This command configures the reaction to an RX LOS.
Parameters	squelch — Specifies to squelch (turn off) the transmit signal on RX LOS.

rx-los-thresh

Syntax	rx-los-thresh <i>threshold</i>
Context	config>port>dwdm>coherent
Description	This command configures the average input power LOS (Loss of Signal) threshold.
Default	-23
Parameters	<i>threshold</i> — Specifies the port's rx los threshold. Values -24.00 to -13.00

sweep

Syntax	sweep start <i>dispersion-start</i> end <i>dispersion-end</i>
Context	config>port>dwdm>coherent
Description	This command allows users to configure the dispersion sweep 'start' and 'end' values for the automatic mode of coherent control. If the user knows the approximate or theoretical residual dispersion of the link, this command can be used to limit the range of sweeping for the automatic control mode and thus achieve faster link up.
Parameters	<i>dispersion-start</i> — Specifies the lower range limit for the dispersion compensation. Values -50000 to 50000 Default -25500 <i>dispersion-end</i> — Specifies the upper range limit for the dispersion compensation. Values -50000 to 50000 Default 2000

target-power

Syntax	target-power <i>power</i>
Context	config>port>dwdm>coherent
Description	This command configures the target transmit optical power for the port.
Default	target-power 1.00
Parameters	<i>power</i> — Specifies the desired average output power in dBm. Values -20.00 to 3.00

rxdtv-adjust

Syntax	[no] rxdtv-adjust
Context	config>port>dwdm
Description	This command enables you to adjust the optical receive decision threshold voltage (RxDTV).
Default	no rxdtv-adjust

tdcm

Syntax	tdcm
Context	config>port>dwdm
Description	This command configures the Tunable Dispersion Compensation Module parameters.

dispersion

Syntax	dispersion <i>dispersion</i>
Context	config>port>dwdm>tdcm
Description	This command allows users to configure the dispersion compensation for the port when manual mode is selected.
Parameters	<i>dispersion</i> — Specifies the dispersion compensation. Values -1200 to 1200 Default 0

mode

Syntax	mode { automatic manual }
Context	config>port>dwdm>tdcm
Description	This command allows users to configure the dispersion algorithm mode used for the port. Manual mode is used when the user knows the residual dispersion on the link. Automatic mode is used to let the software determine the optimal dispersion compensation required. Automatic mode should be used during service commissioning and when the state if the TDCM control is converged, the user can change to manual mode and configure the dispersion compensation found by the software. Because automatic mode uses a search algorithm that will sweep the entire range of dispersion specified in the sweep command, it can take up to 10 minutes for the link to come up. In manual mode, the link can come up in 2 minutes or less.
Parameters	automatic — Sets to automatic mode. manual — Sets to manual mode.

report-alarms

Syntax	[no] report-alarms [nrdy] [meth] [mtl] [unlck] [tlim] [einv] [com]
Context	config>port>dwdm>tdcm
Description	This command allows users to Enable/disable logging of tdcn alarms on the port.
Default	All alarms are enabled
Parameters	nrdy — Reports Tdcn not ready faults. meth — Reports module case temperature high faults. mtl — Reports module case temperature low faults. unlck — Reports thermal control locked faults. tlim — Reports thermal control temperature limit faults. einv — Reports EEPROM invalid faults. com — Reports Tdcn module communication failure faults.

sweep

Syntax	sweep start <i>dispersion-start</i> end <i>dispersion-end</i>
Context	config>port>dwdm>tdcm

Description	This command allows users to configure the dispersion sweep 'start' and 'end' values for the automatic mode of TDCM control. If the user knows the approximate or theoretical residual dispersion of the link, this command can be used to limit the range of sweeping for the automatic control mode and thus achieve faster link up.
Parameters	<i>dispersion-start</i> — Specifies the lower range limit for the dispersion compensation. Values -1200 to 1200 Default -1200 <i>dispersion-end</i> — Specifies the upper range limit for the dispersion compensation. Values -1200 to 1200 Default 1200

wavetracker

Syntax	wavetracker
Context	config>port>dwdm
Description	This command validates whether or not the port supports Wavetracker.

encode

Syntax	encode key1 wave-key key2 wave-key no encode
Context	config>port>dwdm>wavetracker
Description	This command specifies whether or not Wavetracker keys should be encoded on the transmitted optical signal.
Default	no encode
Parameters	<i>wave-key</i> — The <i>wave-key</i> values must be selected based on the currently configured DWDM ITU channel. Both keys must be odd or both keys must be even. One even key and one odd key cannot be configured. The ranges of values for each key are defined in Table 45 .

Table 45 Value Ranges for DWDM ITU Channel

DWDM ITU Channel Number	Key 1 Minimum	Key 1 Maximum	Key 2 Minimum	Key 2 Maximum
17	1276	1290	1760	1774
18	1259	1273	1743	1757
19	1242	1256	1726	1740
20	1225	1239	1709	1723
21	528	542	1072	1086
22	511	525	1055	1069
23	494	508	1038	1052
24	477	491	1021	1035
25	1208	1222	1692	1706
26	460	474	1004	1018
27	443	457	987	1001
28	426	440	970	984
29	409	423	953	967
30	1191	1205	1675	1689
31	392	406	936	950
32	375	389	919	933
33	358	372	902	916
34	341	355	885	899
35	1174	1188	1658	1672
36	324	338	868	882
37	307	321	851	865
38	290	304	834	848
39	273	287	817	831
40	1157	1171	1641	1655
41	256	270	800	814
42	239	253	783	797

Table 45 Value Ranges for DWDM ITU Channel (Continued)

DWDM ITU Channel Number	Key 1 Minimum	Key 1 Maximum	Key 2 Minimum	Key 2 Maximum
17	1276	1290	1760	1774
18	1259	1273	1743	1757
43	222	236	766	780
44	205	219	749	763
45	1140	1154	1624	1638
46	188	202	732	746
47	171	185	715	729
48	154	168	698	712
49	137	151	681	698
50	1123	1137	1607	1621
51	120	134	664	678
52	103	117	647	661
53	86	100	630	644
54	69	83	613	627
55	1106	1120	1590	1604
56	52	66	596	610
57	35	49	579	593
58	18	32	562	576
59	1	15	545	559
60	1089	1103	1573	1587
61	1548	1548	2032	2032
175	3553	3567	4065	4079
185	3536	3550	4048	4062
195	3519	3533	4031	4045
205	3502	3516	4014	4028
215	3840	3854	2304	2318
225	3823	3837	2287	2301

Table 45 Value Ranges for DWDM ITU Channel (Continued)

DWDM ITU Channel Number	Key 1 Minimum	Key 1 Maximum	Key 2 Minimum	Key 2 Maximum
17	1276	1290	1760	1774
18	1259	1273	1743	1757
235	3806	3820	2270	2284
245	3789	3803	2253	2267
255	3485	3499	3997	4011
265	3772	3786	2236	2250
275	3755	3769	2219	2233
285	3738	3752	2202	2216
295	3721	3735	2185	2199
305	3468	3482	3980	3994
315	3704	3718	2168	2182
325	3687	3701	2151	2165
335	3670	3684	2134	2148
345	3653	3667	2117	2131
355	3451	3465	3963	3977
365	3636	3650	2100	2114
375	3619	3633	2083	2097
385	3602	3616	2066	2080
395	3585	3599	2049	2063
405	3434	3448	3946	3960
415	1548	1562	2032	2046
425	1531	1545	2015	2029
435	1514	1528	1998	2012
445	1497	1511	1981	1995
455	3908	3922	2372	2386
465	1480	1494	1964	1978
475	1463	1477	1947	1961

Table 45 Value Ranges for DWDM ITU Channel (Continued)

DWDM ITU Channel Number	Key 1 Minimum	Key 1 Maximum	Key 2 Minimum	Key 2 Maximum
17	1276	1290	1760	1774
18	1259	1273	1743	1757
485	1446	1460	1930	1944
495	1429	1443	1913	1927
505	3891	3905	2355	2369
515	1412	1426	1896	1910
525	1395	1409	1879	1893
535	1378	1392	1862	1876
545	1361	1375	1845	1859
555	3874	3888	2338	2352
565	1344	1358	1828	1842
575	1327	1341	1811	1825
585	1310	1324	1794	1808
595	1293	1307	1777	1791
605	3857	3871	2321	2335

power-control

Syntax [no] power-control**Context** config>port>dwdm>wavetracker>power-control**Description** This command specifies whether the power control loop should be turned on to actively control the laser's launch power to the specified target power. When power-control is disabled, the launch power is set to the laser's maximum achievable power.**Default** no power-control

target-power

Syntax target-power dBm**Context** config>port>dwdm>wavetracker>power-control

Description	This command specifies launch power in dBm for the DWDM Wavetracker-enabled interface.
Default	target-power -20.00
Parameters	<i>dBm</i> — Specifies the desired average output power in dBm.
Values	-22.00 to 3.00

report-alarm

Syntax	[no] report-alarm [encode-fail] [encode-degrade] [power-fail] [power-degrade] [power-high] [power-low] [missing]
Context	config>port>dwdm>wavetracker
Description	This command specifies the alarms which are enabled or outstanding against a Wave Tracker-enabled interface. The no form of the command removes the alarm parameters.
Parameters	encode-fail — Specifies the Encoder failure alarm. encode-degrade — Specifies the Encoder degrade alarm. encode-fail — Specifies the Power control failure alarm. power-degrade — Specifies the Power control degrade alarm. power-high — Specifies the Power control high limit reached alarm. power-low — Specifies the Power control low limit reached alarm. missing — Specifies the wavelength/wavetracker missing alarm.

queue-group

Syntax	queue-group <i>queue-group-name</i> instance <i>instance-id</i> no queue-group
Context	config>port>ethernet>network>egress
Description	This command configures a queue-group instance in the network egress context of a port. Queue-groups containing queues only or policers and queues can be instantiated. When a port is a LAG, one instance of the queue-group is instantiated on each member link. One or more instances of the same queue-group name and/or a different queue-group name can be created in the network egress context of a port. The queue-group-name must be unique within all network egress and access egress queue groups in the system. The queue-group instance-id must be unique within the context of the port.

The **no** version of this command deletes the queue-group instance from the network egress context of the port.

Parameters *queue-group-name* — Specifies the name of the queue group template up to 32 characters.
instance-id — Specifies the identification of a specific instance of the queue-group.
Values 1 to 65535

xgig

Syntax **xgig {lan | wan}**

Context config>port>ethernet

Description This command configures a 10 Gb/s interface to be in Local or Wide Area Network (LAN or WAN) mode. When configuring the port to be in WAN mode, you can change certain SONET/SDH parameters to reflect the SONET/SDH requirements for this port. When you configure a port for LAN mode, all SONET/SDH parameters are pre-determined and not configurable.

Default xgig lan

Parameters **lan** — Sets the port to operate in LAN mode.
wan — Sets the port to operate in WAN mode.

otu

Syntax [**no**] **otu**

Context config>port

Description This command specifies whether or not to enable OTU encapsulation. The port must be shut down before OTU is enabled. This command is valid only for ports on assemblies that support this encapsulation mode. Refer to the appropriate Installation Guide for ports assembly to determine if OTU encapsulation is supported.

Note that OTU cannot be disabled on OTU3 encapsulated OC768 or 40-Gigabit Ethernet by the **no otu** command. Therefore, the default depends on the port type. The default for OTU3 encapsulated OC768 or 40-Gigabit Ethernet is **otu**.

The **no** form of this command disables OTU (clear channel 10GE-LAN/WAN or OC192).

Default no otu

fec

Syntax	fec {g709 enhanced} no fec
Context	config>port>otu>fec
Description	This command enables the Forwarding Error Correction (FEC) encoder/decoder and specifies the FEC encoder/decoder mode to use when enabled. The following rules must be followed: • The port's OTU must be enabled to set or change the FEC mode. • The port must be shut down before changing the FEC mode. • The sf-sd-method must be changed to BIP8 before setting the FEC mode to disabled. Note that FEC cannot be disabled on OTU3 encapsulated OC768 or 40-Gigabit Ethernet by the no fec command. Therefore, the default depends on the port type. The default for OTU3 encapsulated OC768 or 40-Gigabit Ethernet is fec enhanced. The no form of the command disables FEC encoder and decoder.
Default	no fec
Parameters	enhanced — Enables the FEC encoder and decoder with a proprietary enhanced FEC algorithm. g709 — Enables the FEC encoder and decoder with the standard G.709 FEC algorithm.

otu2-lan-data-rate

Syntax	otu2-lan-data-rate {11.049 11.096}
Context	config>port>otu
Description	This command specifies the data rate to use when configured for an OTU encapsulated 10GE-LAN signal. The port must be shut down before changing the 10GE LAN OTU2 data rate.
Default	otu2-lan-data-rate 11.049
Parameters	11.049 — Configures the port to transmit and receive an 11.049 Gb/s synchronous OTU encapsulated 10GE-LAN signal (No fixed stuffing bytes in the OTU2 frame). 11.096 — Configures the port to transmit and receive an 11.096 Gb/s synchronous OTU encapsulated 10GE-LAN signal (with fixed stuffing bytes in the OTU2 frame).

pm-tti

Syntax	pm-tti
Context	config>port>otu
Description	This command enables the context to configure path monitoring trail trace identifier parameters.

expected

Syntax	expected auto-generated expected bytes <i>byte-string</i> [<i>byte-string</i> (up to 64 bytes-strings max, 64 bytes max)] expected string <i>identifier</i> expected use-rx
Context	config>port>otu>pm-tti
Description	This command allows the user to configure the expected RX trail trace identifier (TTI) for path monitoring (PM) in the ODU overhead. This identifier can be a string or a non-printable sequence of bytes. The length of the string or sequence of bytes cannot exceed 64 bytes. This trace should match the far-end port's PM trace. When this trace does not match the received PM trace, the ODU-TIM alarm will be reported if enabled.
Default	Blank (all zeros)
Parameters	auto-generated — Sets the default. <i>identifier</i> — Sets the PM TTI to the string provided by the user. If the string is less than 64 bytes, the remaining bytes will be set to 0. Up to 64 byte strings can be specified in a single statement. <i>byte-string</i> — [byte1 byte2 to byte64]. Sets the PM TTI to the sequence of bytes provided by the user. If the user provides less than 64 bytes, the remaining bytes will be set to 0. use-rx — Copies the received pm-tti to the expected either as a string or a sequence of bytes depending on the received pm-tti data.

mismatch-reaction

Syntax	mismatch-reaction {squelch-rx} no mismatch-reaction
Context	config>port>otu>pm-tti
Description	This command allows the user to configure the consequent action to a pm-tti mismatch. The no form of the command reverts to the default.

Default	n/a, the received traffic is passed through.
Parameters	squelch-rx — Specifies that the received traffic is blocked.

tx

Syntax	tx auto-generated tx bytes <i>byte-string</i> [<i>byte-string...</i> (up to 64 byte-strings max, 64 bytes max)] tx string <i>identifier</i> no tx
Context	config>port>otu>pm-tti
Description	This command enables the user to configure the transmit (tx) trail trace identifier (TTI) for path monitoring (PM) in the ODU overhead. This identifier can be a string or a non-printable sequence of bytes. The length of the string or sequence of bytes cannot exceed 64 bytes. The no form of the command reverts to the default TTI.
Default	Auto-generated in the format of <i>nodename:iomnum/mdanum/portnum/dwdmchan</i> The auto-generated value has five sections: • Nodename — The first section is the name of the node. • iomnum — The second section contains the IOM slot number. • mdanum — The third section contains the MDA slot number. • portnum — The fourth section contains the port number. • dwdmchan — The fifth section contains the DWDM channel number (see DWDM Channel Numbers).
Parameters	auto-generated — Specifies to use the system generated (default) TTI. <i>identifier</i> — Sets the PM TTI to the string provided by the user. If the string is less than 64 bytes, the remaining bytes will be set to 0. <i>byte-string</i> — Sets the PM TTI to the sequence of bytes provided by the user. If the user provides less than 64 bytes, the remaining bytes will be set to 0. A 1 byte sequence of 0xFF will set the default strings. Up to 64 byte strings can be specified in a single statement. Values 0 to FF, in hexadecimal byte notation

psi-payload

Syntax	psi-payload
Context	config>port>otu

Description This command enables the context to configure payload structure identifier payload parameters.

expected

Syntax **expected** *byte*
expected **auto**

Context config>port>otu>psi-payload

Description This command allows the user to configure the expected received payload type value in byte 0 of the Payload structure identifier (PSI) of the OPU overhead. When this values does not match the received value, the OPU-PLM alarm will be reported if it is enabled.

Default 3 for 10GE-LAN/WAN or OC192 with OTU encapsulation; 5 for GFP framed 10GE-LAN with OTU encapsulation.

Parameters **auto** — Sets the expected value to the standard value in the payload type field.
byte — Specifies the expected received payload type value in bytes.

Values [00 to FF] Hexadecimal notation

Default 00

mismatch-reaction

Syntax **mismatch-reaction** {**none** | **squelch-rx**}

Context config>port>otu>psi-payload

Description This command allows the user to configure the consequent action to a psi-payload type mismatch.

Parameters **none** — Specifies the received traffic is passed through.
squelch-rx — Specifies the received traffic is blocked.

sf-sd-method

Syntax **sf-sd-method** {**bip8** | **fec**}

Context config>port>otu>sf-sd-method

Description This command specifies the method used to determine the signal fail and signal degrade alarms. When select the bip8 method is selected, the SM-BIP8 errors are used. When the FEC method is selected, the FEC corrected bits are used.

The following rules must be followed:

- The port's OTU must be enabled to set or change the sf-sd-method.
- The FEC mode must be enhanced or g709 before setting the sf-sd-method to fec.
- The SF threshold must be 5 or higher before setting the sf-sd-method to bip8.

Default	sf-sd-method fec
Parameters	bip8 — The SM-BIP8 errors are used to declare the presence of the Signal Fail and Signal Degrade condition. fec — The FEC corrected bit errors are used to declare the presence of the Signal Fail and Signal Degrade condition.

sf-threshold

Syntax	sf-threshold <i>threshold</i> [coefficient <i>coefficient</i>]
Context	config>port>otu
Description	This command specifies the error rate at which to declare the signal fail condition for the signal fail (SF) threshold. The value represents an error rate of 10E-<value>. The SF threshold must: • Be less than the SD threshold • Be 5 or higher before setting the sf-sd-method to bip8
Default	4
Parameters	<i>threshold</i> — Specifies the signal fail (SF) threshold. Values 3 to 6 Default 5 <i>coefficient</i> — Specifies the coefficient of the SF threshold. Values 10 to 99 Default 10

sf-threshold-clear

Syntax	sf-threshold-clear <i>threshold</i> [coefficient <i>coefficient</i>]
Context	config>port>otu
Description	This command the signal fail (SF) threshold clear.

Parameters	<i>threshold</i> — Specifies the exponent of the error rate, thus an error rate from 10E-3 to 10E-7. Values 5 to 9 Default 6 <i>coefficient</i> — Specifies the coefficient of the SF threshold. Values 10 to 99 Default 10

sd-threshold

Syntax	sd-threshold <i>threshold</i> [coefficient <i>coefficient</i>]
Context	config>port>otu>sd-threshold
Description	This command specifies the error rate at which to declare the signal fail condition for the signal degrade (SD). The value represents an error rate of 10E-<i>value</i>. The SD threshold must be: • greater than the SF threshold. • 5 or higher before setting the sf-sd-method to bip8.
Default	7
Parameters	<i>threshold</i> — Specifies the exponent of the error rate, thus an error rate from 10E-3 to 10E-7. Values 5 to 9 Default 7 <i>coefficient</i> — Specifies the coefficient of the SD threshold. Values 10 to 99 Default 10

sd-threshold-clear

Syntax	sd-threshold-clear <i>threshold</i> [coefficient <i>coefficient</i>]
Context	config>port>otu
Description	This command configures the signal degrade threshold clear.

Parameters	<i>threshold</i> — Specifies the exponent of the error rate, thus an error rate from 10E-3 to 10E-7.
	Values 3 to 10
	Default 8
	<i>coefficient</i> — Specifies the coefficient of the SD threshold.
	Values 10 to 99
	Default 10

sm-tti

Syntax	sm-tti
Context	config>port>otu
Description	This command enables the context to configure section monitoring trail trace identifier parameters.

expected

Syntax	expected auto-generated expected bytes <i>byte-string</i> [<i>byte-string</i> ...(up to 64 byte-strings max, 64 bytes max)] expected string <i>identifier</i> expected use-rx
Context	config>port>otu>sm-tti
Description	This command enables the user to configure the expected RX Trail Trace Identifier (TTI) for Section Monitoring (SM) in the OTU overhead. This identifier can be a string or a non-printable sequence of bytes. The length of the string or sequence of bytes cannot exceed 64 bytes. This trace should match the expected far-end port's SM trace. When this trace does not match the received SM trace, the OTU-TIM alarm will be reported if enabled.
Default	Blank (all zeros)
Parameters	auto-generated — Sets the default. <i>identifier</i> — Sets the PM TTI to the string provided by the user. If the string is less than 64 bytes, the remaining bytes will be set to 0. Up to 64 byte strings can be specified in a single statement. <i>byte-string</i> — [byte1 byte2 to byte64]. Sets the PM TTI to the sequence of bytes provided by the user. If the user provides less than 64 bytes, the remaining bytes will be set to 0. use-rx — Copies the received pm-tti to the expected either as a string or a sequence of bytes depending on the received pm-tti data.

mismatch-reaction

Syntax	mismatch-reaction { none squelch-rx }
Context	config>port>otu>sm-tti
Description	This command allows the user to configure the consequent action to a sm-tti mismatch.
Default	n/a
Parameters	none — Specifies that the received traffic is passed through. squelch-rx — Specifies that the received traffic is blocked.

tx

Syntax	tx <i>byte</i> tx auto				
Context	config>port>otu>psi-payload				
Description	This command allows the user to configure the transmit payload type value in byte 0 of the payload structure identifier (PSI) of the OPU overhead.				
Default	3 for 10GE-LAN/WAN or OC192 with OTU encapsulation; 5 for GFP framed 10GE-LAN with OTU encapsulation.				
Parameters	auto — Transmits the standard value in the payload type field. <i>byte</i> — Specifies the transmit payload type value in bytes. <table> <tr> <td>Values</td><td>[00 to FF] Hexadecimal notation</td></tr> <tr> <td>Default</td><td>00</td></tr> </table>	Values	[00 to FF] Hexadecimal notation	Default	00
Values	[00 to FF] Hexadecimal notation				
Default	00				

tx

Syntax	tx auto-generated tx bytes <i>byte-string</i> [<i>byte-string</i> ...(up to 64 byte-strings max, 64 bytes max)] tx string <i>identifier</i> no tx
Context	config>port>otu>sm-tti
Description	This command allows the user to configure the transmit (tx) trail trace identifier (TTI) for section monitoring (SM) in the OTU overhead. This identifier can be a string or a non-printable sequence of bytes. The length of the string or sequence of bytes cannot exceed 64 bytes. The no form of the command reverts to the default TTI.

Default Auto-generated in the format of *nodename:iomnum/mdanum/portnum/dwdmchan*

The auto-generated value has five sections:

- Nodename — The first section is the name of the node.
- iomnum — The second section contains the IOM slot number.
- mdanum — The third section contains the MDA slot number.
- portnum — The fourth section contains the port number.
- dwdmchan — The fifth section contains the DWDM channel number (see [DWDM Channel Numbers](#)).

Parameters **auto-generated** — Specifies to use the system generated (default) TTI.

identifier — Sets the SM TTI to the string provided by the user. If the string is less than 64 bytes, the remaining bytes will be set to 0. Up to 64 byte strings can be specified in a single statement.

byte-string — Sets the SM TTI to the sequence of bytes provided by the user. If the user provides less than 64 bytes, the remaining bytes will be set to 0. A 1 byte sequence of 0xFF will set the default strings.

Values 0 to FF, in hexadecimal byte notation

psi-tti

Syntax **psi-tti**

Context config>port>otu

Description This command enables the context to configure payload structure identifier trail trace identifier parameters.

tx

Syntax **tx {string *identifier* | bytes *byte-sequence* | auto-generated}**

Context config>port>otu>psi-trace

Description This command allows the user to configure the transmit trace in bytes 1 to 255 (skipping byte 0) of the payload structure identifier (PSI) of the OPU overhead. This identifier can be a string or a non-printable sequence of bytes. The length of the string or sequence of bytes cannot exceed 255 bytes.

Default Blank (all zeros)

Parameters **auto-generated** — Sets the default PSI trace.

identifier — Sets the PSI trace to the string provided by the user. If the string is less than 255 bytes, the remaining bytes will be set to 0.

byte-sequence — [byte1 byte2 to byte64] Sets the PSI trace to the sequence of bytes provided by the user. If the user provides less than 64 bytes, the remaining bytes will be set to 0. A 1 byte sequence of 0xFF will set the default strings.

Values 0 to FF, in hexadecimal byte notation

mismatch-reaction

Syntax	mismatch-reaction {none squelch-rx}
Context	config>port>otu>psi-tti
Description	This command allows the user to configure the consequent action to a psi-tti mismatch.
Default	n/a
Parameters	none — Specifies the received traffic is passed through. squelch-rx — Specifies the received traffic is blocked.

async-mapping

Syntax	[no] async-mapping
Context	config>port>otu
Description	This command allows the user to configure the port to support asynchronous mapping of the payload inside the OTU. If the port is configured for async-mapping and the payload clock is asynchronous to the OTU clock, there will be positive or negative pointer justification that will show up in the OTU statistics and the data will be received error free. If the port is configured for synchronous mapping and the received data is asynchronously mapped, there will be errors in the received data. async-mapping is the only mode of operation that is supported on the OTU3 encapsulated 40-Gigabit Ethernet and therefore the 'no async-mapping' is not supported on that port type and the default on the is async-mapping. The no form of this command configures the port to receive synchronously mapped data.
Default	no async-mapping

report-alarms

Syntax	[no] report-alarms [loc] [los] [lof] [lom] [otu-ais] [otu-ber-sf] [otu-ber-sd] [otu-bdi] [otu-tim] [otu-iae] [otu-biae] [fec-sf] [fec-sd] [fec-fail] [fec-uncorr] [odu-ais] [odu-oci] [odu-ick] [odu-bdi] [odu-tim] [opu-tim] [opu-plm]
Context	config>port>otu

Description This command enables OTU alarms. Specify specific alarms to add to the list of reported alarms.

The **no** form of the command disables OTU alarm reporting.

Default loc, los, lof, lom, otu-ber-sf, otu-bdi, fec-sf

Parameters **alarms** — Refer to [Table 46](#) for alarm descriptions.

Table 46 Alarm Descriptions

Alarm	Description
loc	Loss of lock.
lof	Loss of OTU framing.
lom	Loss of Multi-frame.
los	Loss of signal transitions on the data.
otu-ais	OTU Alarm Indication Signal (all 1s, overwrites all OTU overhead, even framing bytes).
otu-ber-sf	SM Signal Fail (based on BPI8).
otu-ber-sd	SM Signal Degrade (based on BPI8).
otu-bdi	SM Backward defect indication.
otu-tim	SM Trace Id Mismatch.
otu-iae	SM Incoming Alignment Error.
otu-biae	SM Backward Incoming Alignment Error.
fec-sf	Signal Fail (based on FEC corrected bits).
fec-sd	Signal Degrade (based on FEC corrected bits).
fec-fail	FEC Mode mismatch (EFEC-GFEC) or High Uncorrectable rate (>10E-2).
fec-uncorr	One or More Uncorrectable FEC errors.
odu-ais	ODU Alarm Indication Signal.
odu-oci	ODU Open connection Indication.
odu-lck	ODU Locked.
odu-bdi	PM Backward Defect indication.
odu-tim	PM Trace Id Mismatch.
opu-tim	OPU PSI Trace Mismatch.

Table 46 Alarm Descriptions (Continued)

Alarm	Description (Continued)
opu-plm	OPU PSI Payload Type Mismatch.

hybrid-buffer-allocation

Syntax	hybrid-buffer-allocation
Context	config>port
Description	This command enables the context for configuring hybrid port buffer allocation parameters.

egr-weight

Syntax	egr-weight access <i>access-weight</i> network <i>network-weight</i> no egr-weight
Context	config>port>hybrid-buffer-allocation
Description	This command configures the sharing of the egress buffers allocated to a hybrid port among the access and network contexts. By default, it is split equally between network and access. The no form of this command restores the default values for the egress access and network weights.
Parameters	<i>access-weight</i> — Specifies the access weight as an integer. Values 0 to 100 Default 50 <i>network-weight</i> — Specifies the network weight as an integer. Values 0 to 100 Default 50

ing-weight

Syntax	ing-weight access <i>access-weight</i> network <i>network-weight</i> no ing-weight
Context	config>port>hybrid-buffer-allocation
Description	This command configures the sharing of the ingress buffers allocated to a hybrid port among the access and network contexts. By default, it is split equally between network and access.

The **no** form of this command restores the default values for the ingress access and network weights.

Parameters	<i>access-weight</i> — Specifies the access weight as an integer.
	Values 0 to 100
	Default 50
	<i>network-weight</i> — Specifies the network weight as an integer.
	Values 0 to 100
	Default 50

modify-buffer-allocation-rate

Syntax	modify-buffer-allocation-rate
Context	config>port
Description	This command enables the context to configure ingress and egress percentage of rate parameters. This command only applies to physical ports (for example, it will not work on APS or similar logical ports). The percentage of rate commands are used to define a percentage value that affects the amount of buffers used by ingress and egress port managed buffer space. Enter the modify-buffer-allocation-rate context when editing the port's percentage of rate commands.

egr-percentage-of-rate

Syntax	egr-percentage-of-rate <i>egr-rate-percentage</i> no egr-percentage-of-rate
Context	config>port>modify-buffer-allocation-rate
Description	The egr-percentage-of-rate command increases or decreases the active bandwidth associated with the egress port that affects the amount of egress buffer space managed by the port. Changing a ports active bandwidth using the egr-percentage-of-rate command is an effective means of artificially lowering the buffers managed by one egress port and giving them to other egress ports on the same MDA. The egr-percentage-of-rate command accepts a percentage value that increases or decreases the active bandwidth based on the defined percentage. A value of 50% causes the active bandwidth to be reduced by 50%. A value of 150% causes the active bandwidth to be increased by 50%. Values from 1 to 1000 percent are supported. A value of 100 (the default value) is equivalent to executing the no egr-percentage-of-rate command and restores the egress active rate to the normal value.

The **no** form of the command removes any artificial increase or decrease of the egress active bandwidth used for egress buffer space allocation to the port. The **no egr-percentage-of-rate** command sets rate-percentage to 100%.

- Parameters** *egr-rate-percentage* — The rate-percentage parameter is required and defines the percentage value used to modify the current egress active bandwidth of the port. This does not actually change the bandwidth available on the port in any way. The defined rate-percentage is multiplied by the egress active bandwidth of the port. A value of 150 results in an increase of 50% (1.5 x Rate).
- Values** 1 to 1000
- Default** 100 (no change to active rate)

ing-percentage-of-rate

- Syntax** **ing-percentage-of-rate** *ing-rate-percentage*
no ing-percentage-of-rate
- Context** config>port>modify-buffer-allocation-rate
- Description** This command increases or decreases the active bandwidth associated with the ingress port that affects the amount of ingress buffer space managed by the port. Changing a port's active bandwidth using the **ing-percentage-of-rate** command is an effective means of artificially lowering the buffers managed by one ingress port and giving them to other ingress ports on the same MDA.
- The **ing-percentage-of-rate** command accepts a percentage value that increases or decreases the active bandwidth based on the defined percentage. A value of 50% causes the active bandwidth to be reduced by 50%. A value of 150% causes the active bandwidth to be increased by 50%. Values from 1 to 1000 percent are supported.
- A value of 100 (the default value) is equivalent to executing the **no ing-percentage-of-rate** command and restores the ingress active rate to the normal value.
- The **no** form command removes any artificial increase or decrease of the ingress active bandwidth used for ingress buffer space allocation to the port. The **no ing-percentage-of-rate** command sets rate-percentage to 100%.
- Parameters** *ing-rate-percentage* — The rate-percentage parameter is required and defines the percentage value used to modify the current ingress active bandwidth of the port. This does not actually change the bandwidth available on the port in any way. The defined rate-percentage is multiplied by the ingress active bandwidth of the port. A value of 150 results in an increase of 50% (1.5 x Rate).
- Values** 1 to 1000
- Default** 100 (no change to active rate)

monitor-agg-egress-queue-stats

Syntax	[no] monitor-agg-egress-queue-stats
Context	config>port
Description	This command enables the monitoring of aggregate egress queue statistics on the port. All queues on the port are monitored, including SAP egress, network egress, subscriber egress, and egress queue group queues, as well as system queues that can be used, for example, to send port-related protocol packets (LACP, EFM, and so on). The aggregate in-profile, out-of-profile, and total statistics are provided for both forwarded and dropped packets and octets. Monitoring of aggregate statistics is supported on PXC sub-ports but not on a PXC physical port. It is also not supported on satellite ports or ports on an HSMDA. The no form of the command disables aggregate egress queue statistics monitoring on the specified port.

egress-scheduler-override

Syntax	egress-scheduler-override [create] no egress-scheduler-override
Context	config>port>sonet-sdh>path config>port>ethernet config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command applies egress scheduler overrides. When a port scheduler is associated with an egress port, it is possible to override the following parameters: • The max-rate allowed for the scheduler. • The maximum rate for each priority level 8 through 1. • The CIR associated with each priority level 8 through 1. See the <i>7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide</i> for command syntax and usage for the port-scheduler-policy command. The no form of this command removes all override parameters from the egress port or channel scheduler context. Once removed, the port scheduler reverts all rate parameters back to the parameters defined on the port-scheduler-policy associated with the port.
Parameters	create — Mandatory while creating an entry.

level

Syntax	level <i>priority-level</i> rate <i>pir-rate</i> [cir <i>cir-rate</i>] level <i>priority-level</i> percent-rate <i>pir-percent</i> [percent-cir <i>cir-percent</i>] no level <i>priority-level</i>
Context	config>port>ethernet>egress-scheduler-override config>port>sonet-sdh>path>egress-scheduler-override config>port>tdm>ds1>channel-group>egress-scheduler-override config>port>tdm>ds3>egress-scheduler-override config>port>tdm>e1>egress-scheduler-override config>port>tdm>e3>egress-scheduler-override
Description	This command overrides the maximum and CIR rate parameters for a specific priority level on the port or channel's port scheduler instance. When the level command is executed for a priority level, the corresponding priority level command in the port-scheduler-policy associated with the port is ignored. The override level command supports the keyword max for the rate and cir parameter. When executing the level override command, at least the rate or cir keywords and associated parameters must be specified for the command to succeed. The no form of this command removes the local port priority level rate overrides. Once removed, the port priority level will use the port scheduler policies level command for that priority level.
Parameters	<i>priority-level</i> — Identifies which of the eight port priority levels are being overridden. Values 1 to 8 <i>pir-rate</i> — Overrides the port scheduler policy's maximum level rate and requires either the max keyword or a rate defined in kilobits per second to follow. Values max, 1 to 3200000000 kb/s <i>cir-rate</i> — Overrides the port scheduler policy's within-cir level rate and requires either the max keyword or a rate defined in kilobits per second to follow. Values max, 0 to 3200000000 kb/s <i>pir-percent</i> — Specifies the PIR as a percentage. Values 0.01 to 100.00 <i>cir-percent</i> — Specifies the CIR as a percentage. Values 0.00 to 100.00 max — removes any existing rate limit imposed by the port scheduler policy for the priority level allowing it to use as much total bandwidth as possible.

max-rate

Syntax **max-rate** *rate*

	max-rate <i>percent-rate</i> no max-rate
Context	config>port>ethernet>egress-scheduler-override>level>rate config>port>ethernet>egress-scheduler-override config>port>sonet-sdh>path>egress-scheduler-override config>port>tdm>ds1>channel-group>egress-scheduler-override config>port>tdm>ds3>egress-scheduler-override config>port>tdm>e1>channel-group>egress-scheduler-override config>port>tdm>e3>egress-scheduler-override
Description	This command overrides the max-rate parameter found in the port-scheduler-policy associated with the port. When a max-rate is defined at the port or channel level, the port scheduler policies max-rate parameter is ignored. The egress-scheduler-override max-rate command supports a parameter that allows the override command to restore the default of not having a rate limit on the port scheduler. This is helpful when the port scheduler policy has an explicit maximum rate defined and it is desirable to remove this limit at the port instance. The no form of this command removes the maximum rate override from the egress port or channels port scheduler context. Once removed, the max-rate parameter from the port scheduler policy associated with the port or channel will be used by the local scheduler context.
Parameters	<i>rate</i> — Specifies the explicit maximum frame based bandwidth limit. This value overrides the QoS scheduler policy rate. Values 1 to 3200000000, max <i>percent-rate</i> — Values 0.01 to 100.00

egress-scheduler-policy

Syntax	egress-scheduler-policy <i>port-scheduler-policy-name</i> no egress-scheduler-policy
Context	config>port>ethernet config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command enables the provisioning of an existing port-scheduler-policy to a port or channel. The egress-scheduler-override node allows for the definition of the scheduler overrides for a specific port or channel.

When a port scheduler is active on a port or channel, all queues and intermediate service schedulers on the port are subject to receiving bandwidth from the scheduler. Any policers, queues, or schedulers with port-parent associations are mapped to the appropriate port priority levels based on the port-parent command parameters. Any policers, queues, or schedulers that do not have a port-parent or valid intermediate scheduler parent defined are treated as orphaned and are handled based on the port scheduler policies default or explicit orphan behavior.

The port scheduler maximum rate and priority level rate parameters may be overridden to allow unique values separate from the port-scheduler-policy-name attached to the port or channel. Use the **egress-scheduler-override** command to specify the port or channel specific scheduling parameters.

The command used to associate an egress scheduler policy on the port is also used for the HSMDB. HSMDB policies should be associated with HSMDB ports.

The **no** form of this command removes a port scheduler policy from an egress port or channel. Once the scheduler policy is removed, all orphaned policers, queues, and schedulers revert to a free running state governed only by the local queue or scheduler parameters. This includes any queues or schedulers with a port-parent association.

Parameters *port-scheduler-policy-name* — Specifies an existing port-scheduler-policy configured in the **config>qos** context. The name can be up to 32 characters.

elmi

Syntax **elmi**

Context config>port>ethernet

Description This command configures Ethernet Local Management Interface (E-LMI) parameters for the Ethernet port. E-LMI is only supported on Ethernet access ports with Dot1q encapsulation type.

mode

Syntax **mode {none | uni-n}**

Context config>port>ethernet>elmi

Description This command configures the Ethernet LMI mode.

Parameters **none** — Specifies that the E LMI mode is set to none.
uni-n — Specifies that the E LMI mode is set to UNI-N.

n393

Syntax	n393 [<i>value</i>] no n393
Context	config>port>ethernet>elmi
Description	This command configures the monitored count of consecutive errors.
Parameters	<i>value</i> — Specifies the monitored count of consecutive errors. Values 2 to 10

t391

Syntax	t391 [<i>value</i>] no t391
Context	config>port>ethernet>elmi
Description	This command configures the polling timer for UNI-C.
Parameters	<i>value</i> — Specifies the polling timer for UNI-C. Values 5 to 30

t392

Syntax	t392 [<i>value</i>] no t392
Context	config>port>ethernet>elmi
Description	This command configures the polling verification timer for UNI-N.
Parameters	<i>value</i> — Specifies the polling verification timer for UNI-N. Values 5 to 30

mode

Syntax	mode { access network hybrid } no mode
Context	config>lag config>port>ethernet config>port>sonet-sdh>path config>port>tdm>ds1>channel-group

```
config>port>tdm>ds3  
config>port>tdm>e1>channel-group  
config>port>tdm>e3
```

Description This command configures an Ethernet port, TDM channel, or SONET/SDH path (sub-port) for **access, network or hybrid** mode operation.

An **access** port or channel is used for customer facing traffic on which services are configured. A Service Access Point (SAP) can only be configured on an access port or channel. When a port is configured for access mode, the appropriate **encap-type** must be specified to distinguish the services on the port or SONET path. Once an Ethernet port, a TDM channel or a SONET path has been configured for access mode, multiple services can be configured on the Ethernet port, a TDM channel or SONET path. Note that ATM, Frame Relay, and cHDLC port parameters can only be configured in the access mode.

An access port or channel is used for customer facing traffic on which services are configured. A Service Access Point (SAP) can only be configured on an access port or channel. When a port is configured for access mode, the appropriate encap-type must be specified to distinguish the services on the port or SONET path. Once an Ethernet port, a TDM channel or a SONET path has been configured for access mode, multiple services can be configured on the Ethernet port, a TDM channel or SONET path. Note that ATM, Frame Relay, and cHDLC port parameters can only be configured in the access mode.

A network port or channel participates in the service provider transport or infrastructure network when a network mode is selected. When the network option is configured, the encap-type cannot be configured for the port/channel.

When network mode is selected on a SONET/SDH path, the appropriate control protocols are activated when the need arises. For example, configuring an IP interface on the SONET path activates IPCP while the removal of the IP interface causes the IPCP to be removed. The same applies for MPLS, MPLSCP, and OSICP. When configuring a SONET/SDH port, the mode command must be entered in the channel context or an error message is generated.

A hybrid Ethernet port allows the combination of network and access modes of operation on a per-VLAN basis and must be configured as either dot1q or QinQ encapsulation.

When the hybrid port is configured to the dot1q encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and an unused VLAN tag value. The format is *<port-id>:qtag1*. A SAP of format *<port-id>:** also supported.

The user configures a network IP interface under **config>router>if>port** by providing the port name which consists of the port-id of the hybrid mode port and an unused VLAN tag value. The format is *<port-id>:qtag1*. The user must explicitly enter a valid value for qtag1. The *<port-id>:** value is not supported on a network IP interface. The 4096 VLAN tag space on the port is shared among VLAN SAPs and VLAN network IP interfaces.

When the hybrid port is configured to QinQ encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and the outer and inner VLAN tag values. The format is <port-id>:qtag1.qtag2. A SAP of format <port-id>: qtag1.* is also supported. The outer VLAN tag value must not have been used to create an IP network interface on this port. In addition, the qtag1.qtag2 value combination must not have been used by another SAP on this port.

The user configures a network IP interface under **config>router>if>port** by providing the port name which consists of the port-id of the hybrid mode port and a VLAN tag value. The format is <port-id>:qtag1.*. An outer VLAN tag qtag2 of * creates an IP network interface. In addition, the qtag1.qtag2 value combination must not have been used on another SAP or IP network interface on this port.

The **no** form of this command restores the default.

Default	mode network — Configures the Ethernet port, TDM channel or SONET path for transport network use. mode access — Default channel/port mode for channelized, ASAP, and ATM MDAs.
Special Cases	SONET/SDH Path — When network mode is selected, the appropriate control protocols are activated when the need arises. For example, configuring an IP interface on the SONET path activates IPCP while the removal of the IP interface causes the IPCP to be removed. The same applies for MPLS, MPLSCP, and OSICP. When configuring a SONET/SDH port, the mode command must be entered in the channel context or an error message is generated.
Parameters	network — Configures the Ethernet port, TDM channel or SONET path as service access. access — Configures the Ethernet port, TDM channel or SONET path for transport network use. hybrid — Configures the Ethernet port for hybrid use.

mac

Syntax	mac <i>ieee-address</i> no mac
Context	config>eth-tunnel config>lag config>port>ethernet config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3

Description	This command assigns a specific MAC address to an Ethernet port, Link Aggregation Group (LAG), Ethernet tunnel, or BCP-enabled port or sub-port. Only one MAC address can be assigned to a port. When multiple mac commands are entered, the last command overwrites the previous command. When the command is issued while the port is operational, IP will issue an ARP, if appropriate, and BPDUs are sent with the new MAC address. The no form of this command returns the MAC address to the default value.
Default	A default MAC address is assigned by the system from the chassis MAC address pool.
Parameters	<i>ieee-address</i> — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

mtu

Syntax	mtu <i>mtu-bytes</i> no mtu
Context	config>port>ethernet config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command configures the maximum payload MTU size for an Ethernet port, PPP-enabled port or sub-port and Frame Relay-enabled port or subport. The Ethernet port level MTU parameter indirectly defines the largest physical packet the port can transmit or the far-end Ethernet port can receive. Packets that cannot be fragmented at egress and exceed the MTU are discarded. The value specified for the MTU includes the destination MAC address, source MAC address, the Ethertype or Length field and the complete Ethernet payload. The MTU value does not include the preamble, start of frame delimiter or the trailing CRC. PoS channels use the MTU to define the largest PPP payload a PoS frame may contain. A significant difference between SONET/SDH PoS channel and Ethernet physical MTU values is the overhead considered part of the framing method and the overhead considered to be part of the application using the frame. In Ethernet, the preamble, start of frame delimiter and the CRC are considered part of the framing overhead and not part of the frame payload. For a PoS channel, the HDLC framing overhead is not included in the physical MTU; only the PPP and PPP payload are included. If the port mode or encapsulation type is changed, the MTU assumes the default values of the new mode or encapsulation type. The no form of this command restores the default values.

Default The default MTU value depends on the (sub-)port type, mode and encapsulation and are listed in [Table 47](#):

Table 47 Default MTU Values

Type	Mode	Encap Type	Default (Bytes)
10/100, Gig, or 10GigE	Access	null	1514
10/100, Gig, or 10GigE	Access	dot1q	1518
10/100, Gig, or 10GigE	Access	q-in-q	1522
SONET/SDH or TDM	Access	mpls	1506
SONET/SDH or TDM	Access	bcp-null	1518
SONET/SDH or TDM	Access	bcp-dot1q	1522
SONET/SDH or TDM	Access	ipcp	1502
SONET/SDH or TDM	Access	frame-relay	1578
ATM, SONET/SDH or TDM	Access	atm	1524
10/100 or 100FX Ethernet	Network	null	1514
10/100 or 100FX Ethernet	Network	dot1q	1518
SONET/SDH	Network	ppp-auto	1524

Parameters *mtu-bytes* — Sets the maximum allowable size of the MTU, expressed as an integer.

Values

512 to 9212	config>port>ethernet
512 to 9208	config>port>sonet-sdh>path
512 to 9208	config>port>tdm>ds1>channel-group
512 to 9208	config>port>tdm>ds3
512 to 9208	config>port>tdm>e1>channel-group
512 to 9208	config>port>tdm>e3

network

Syntax **network**

Context config>port>tdm>ds1>channel-group
config>port>tdm>e1>channel-group

Description This command enables the context to configure network channel group parameters.

queue-policy

Syntax	queue-policy <i>name</i> no queue-policy
Context	config>port>tdm>ds1>channel-group>network config>port>tdm>e1>channel-group>network
Description	This command specifies an existing network policy to apply to the channel group.
Parameters	<i>name</i> — Specifies an existing network-queue policy name. The name can be up to 32 characters.

queue-policy

Syntax	queue-policy <i>name</i> no queue-policy
Context	config>card>mda>network>ingress config>port>sonet-sdh>path>network
Description	This command specifies the network-queue policy which defines queue parameters such as CBS, high priority only burst size, MBS, CIR and PIR rates, as well as forwarding-class to queue mappings. The network-queue policy is defined in the config>qos>network-queue context.
Default	queue-policy default
Parameters	<i>name</i> — Specifies an existing network-queue policy name. The name can be up to 32 characters long.

ppp

Syntax	ppp
Context	config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command enables access to the context to configure the LCP operational parameters for a SONET/SDH PoS link, a DS--3/E-3 port or channel, a DS-1/E-1 channel or a DS-0 channel.

compress

Syntax	compress {acfc [pfc] pfc [acfc]} no compress
Context	config>port>tdm>ds1>channel-group>ppp config>port>tdm>e1>channel-group>ppp
Description	This command enables and disables Protocol Field Compression (PFC) per RFC 1661, <i>The Point-to-Point Protocol (PPP)</i> , Section 6.5 and Address and Control Field Compression (ACFC) as per Section 6.6. This command is only supported on DS-1 and E-1 channel groups on ASAP MDAs. The no form of the command disables the header compression.
Default	no compress
Parameters	acfc — Specifies address and control field compression. pfc — Specifies protocol field compression.

ber-sf-link-down

Syntax	[no] ber-sf-link-down
Context	config>port>tdm>ds1>channel-group>ppp config>port>tdm>e1>channel-group>ppp
Description	This command enables the port down on BER-SF alarm. When enabled, the link will be placed out of service once ber-sf is detected. The no form of the command reverts to normal operation where the link remains in-service when ber-sf is encountered.
Default	no ber-sf-link-down

report-alarm

Syntax	[no] report-alarm [ais] [los] [oof] [rai] [looped]
Context	config>port>tdm> ds3 config>port>tdm> e3
Description	This command enables logging of DS-3 and E-3 alarms for a DS-3/E-3 port or channel. The no form of this command disables logging of the specified alarms.

Parameters	ais — Reports alarm indication signal errors. When configured, ais alarms are not raised and cleared. Default ais alarms are issued los — Reports loss of signal errors. When configured, los traps are not raised and cleared. Default los traps are issued oof — Reports out-of-frame errors. When configured, oof alarms are not raised and cleared. Default oof alarms are not issued rai — Reports resource availability indicator events. When configured, rai events are not raised and cleared. Default rai alarms are not issued looped — Reports looped packets errors. Default looped alarms are not issued

scramble

Syntax	[no] scramble
Context	<pre>config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3</pre>
Description	This command enables payload scrambling on channel groups. Scrambling randomizes the pattern of 1s and 0s carried in a SONET frame. Rearranging or scrambling the pattern prevents continuous strings of all 1s or all 0s and meets the needs of physical layer protocols that rely on sufficient transitions between 1s and 0s to maintain clocking. For ATM, this command enables or disables ATM cell-level payload scrambling/descrambling using $x^{43}+1$ polynomial as defined in ITU-T I.432.1. Scrambling is enabled by default for the ATM path/channel. Note that this scrambling is done in addition to SONET/SDH frame scrambling/descrambling, which is always enabled in the framer. The no form of this command disables scrambling.
Default	no scramble

keepalive

Syntax	keepalive <i>time-interval</i> [dropcount <i>drop-count</i>]

no keepalive

Context	<pre>config>port>sonet-sdh>path>ppp config>port>tdm>ds1>channel-group>cisco-hdlc config>port>tdm>ds1>channel-group>ppp config>port>tdm>ds3>ppp config>port>tdm>e1>channel-group>ppp config>port>tdm>e3>ppp</pre>
Description	This command sets the keepalive interval. The no form of this command returns the interval to the default value.
Default	keepalive 10
Parameters	<i>time-interval</i> — Specifies the time, in seconds, between keepalive messages, expressed as a decimal integer. Values 1 to 60 <i>drop-count</i> — Specifies the number of consecutive keepalive failed request attempts or remote replies that can be missed after which the port is operationally downed. Values 1 to 255

2.20.2.11 Port XC Commands

The port xc commands are supported on the 7450 ESS only.

port-xc

Syntax	port-xc
Context	config
Description	This command enables the context to configure port-cross connect functionality.

pxc

Syntax	<pre>pxc pxc-id [create] no pxc pxc-id</pre>
Context	config>port-xc

Description	This command creates a port cross-connect (PXC) object. Referencing an Ethernet port within the PXC object will automatically configure this Ethernet port as a loopback port. The node will automatically create two PXC sub-ports under this Ethernet port. The configuration of PXC sub-ports can be accessed through the CLI.
Parameters	<i>pxc-id</i> — Specifies the port cross-connect identifier.
Values	1 to 64

port

Syntax	port <i>port-id</i> no port
Context	config>port-xc>pxc
Description	This command configures the referenced Ethernet port as a loopback or a cross-connect port (PXC). Once this command is executed, the system automatically creates two PXC sub-ports under this Ethernet port. The two PXC sub-ports are logical configurations used by the node to transmit traffic bi-directionally through a single physical port that is internally cross-connected. The physical PXC port does not require any external connectivity or optical transceivers to function properly. Consequently, all optic-related alarms are disabled on the port. The physical PXC port is automatically configured as a hybrid port. The MTU is preset to 9212 bytes, The encapsulation type is set to dot1q and dot1x tunneling is turned on. Since the PXC is using a single physical port to transmit traffic in both directions, the nominal port bandwidth is asymmetrically divided between the two directions. For example, a 10Gb/s Ethernet port in PXC mode can accommodate nine Gb/s of traffic in one direction and one Gb/s in the other. Any other ratio can be achieved as long as the sum of the bandwidth of the two PXC sub-ports does not exceed the bandwidth capacity of the physical port (10 Gb/s in this case). The following apply to PXC ports: • Only unused physical ports (not associated with an interface or SAP) can be referenced inside of a PXC ID configuration. • Once inside of a PXC ID configuration, the physical port cannot be removed from that PXC ID configuration if the corresponding PXC sub-ports are currently in use. • Once inside of a PXC ID configuration, the physical port cannot be used outside of the PXC context. For example, a regular IP interface cannot use this physical port, or a SAP on that port cannot be associated with a service. • A physical port can be associated with only one PXC ID configuration.
Parameters	<i>port-id</i> — Specifies the physical port in the <i>slot/mda/port</i> format.

pxc-pxc-id.sub-port-id

Syntax	pxc-pxc-id.sub-port-id
Context	config>port
Description	This command enables access to PXC sub-port level parameters. The PXC sub-ports are automatically created once the external Ethernet port is configured inside of an PXC object. The PXC sub-ports are by default administratively disabled (shutdown). In order for PXC sub-ports to become operational, both, the underlying external Ethernet port and the PXC object must be operationally up.
Parameters	<i>pxc-id</i> — Specifies the unique identifier of this PXC. Values 1 to 64 <i>sub-port-id</i> — When this the <i>pxc-id</i> is configured, two logical sub-ports are automatically created. These logical sub-ports are used to create two paths within the loop; one upstream path, and one downstream path. These sub-ports are destroyed when either this PXC row is destroyed, this object is de-provisioned. Values a, b

2.20.2.12 Forwarding Path Extension (FPE) Commands

fwd-path-ext

Syntax	fwd-path-ext
Context	config
Description	This command provides context for configuring Forwarding Path Extensions (FPE). FPE is utilized by certain applications that rely on PXC functionality. Its purpose is to simplify configuration of such applications.

fpe

Syntax	fpe fpe-id [create] no fpe fpe-id
Context	config>fwd-path-ext
Description	This command configures an FPE object which associates the application with a PXC (paired set of PXC sub-ports or a paired set of PXC based LAGs). The no form of the command disables the FPE object association.

Parameters *fpe-id* — Specifies the FPE ID.

Values 1 to 64

create — Keyword used to associate the queue group. The create keyword requirement can be enabled or disabled in the **environment>create** context.

path

Syntax **path xc-a lag-id xc-b lag-id**
path pxc pxc-id
no path

Context config>fwd-path-ext>fpe

Description This command references a PXC (pair of PXC sub-ports) and consequently create an association between the PXC and the application which is referenced under the same FPE object. Each application will utilize the PXC in the form of an internal cross-connect. The exact use and internal provisioning of this cross-connect depends on the application itself.

The **no** form of the command removes the reference and association from the configuration.

Default no path

Parameters **xc-a lag-id** — Specifies the LAG identifier associated with one side of the cross-connect. The operator has the freedom to associate **xc-a** with LAG ID containing either sub-ports.a or sub-ports.b. In other words, the system does not perform automatic check that will ensure a match between **xc-a** and the LAG ID containing sub-ports.a.

Values 1 to 800

xc-b lag-id — Specifies the LAG identifier associated with one side of the cross-connect. The operator has the freedom to associate **xc-a** with LAG ID containing either sub-ports.a or sub-ports.b.

Values 1 to 800

pxc-id — Specifies the PXC identifier, the PXC construct that contains a physical port in a loopback mode that provides the cross-connect capability. The system creates two paired sub-ports on top of this physical port and each of these two sub-ports forwards traffic in one direction over the loopback. One sub-port is associated with the transit side of the loopback, while the other sub-port is associated with the termination side (see PXC Configuration Guides for further explanation).

Values 1 to 64

pw-port

Syntax **[no] pw-port**

Context config>fwd-path-ext>fpe

Description This command informs the system about the type of the cross-connect that is required in order to terminate an external tunnel to an anchored PW port. The system automatically builds the internal infrastructure required to perform the tunnel termination on a PW port.

PW ports support the following types of tunnels:

- GRE/MPLS PW
with SDP of type MPLS or GRE
- L2oGRE
bridged Ethernet over GRE, where GRE protocol number is 0x6558

The **no** form of the command removes the cross-connect type from the configuration.

Default no pw-port

sub-mgmt-extensions

Syntax **[no] sub-mgmt-extensions**

Context config>fwd-path-ext>fpe

Description This command configures FPE for subscriber management extensions. The FPE cannot be used for other applications but can be used for multiple subscriber management applications.

The **no** version of this command disables FPE for subscriber management extensions.

Default no sub-mgmt-extensions

vxlan-termination

Syntax **vxlan-termination [router *router-name*]**
vxlan-termination service-name *service-name*
no vxlan-termination

Context config>fwd-path-ext>fpe

Description This command informs the system about the cross-connect type that is required for non-system IPv4 and IPv6 VXLAN termination. Internally, it triggers the automatic creation of two internal IP interfaces in the PXC ports and enables those internal interfaces to process and terminate VXLAN.

If no parameters are used, the VXLAN termination occurs in the base router; however, when the FPE is used for static VXLAN termination (no BGP-EVPN services), non-system IPv4 and IPv6 VXLAN can be terminated in a VPRN service. In this case, the VPRN router instance or service name must be configured with the **vxlan-termination** command.

The **no** form of the command disables the cross-connect type from the configuration.

Default no vxlan-termination

Parameters *router-name* — Specifies the router instance for VXLAN termination.

Values

router-name: *router-name* or *vprn-svc-id*

<i>router-name</i>	“Base”
<i>vprn-svc-id</i>	1 to 2147483647

Default Base

service-name — Specifies the service name that identifies the VPRN for VXLAN termination, up to 64 characters.

sdp-id-range

Syntax **sdp-id-range** from *id* to *id*
no sdp-id-range

Context config>fwd-path-ext

Description This command reserves an SDP ID range used by the FPE based PW-Port and VXLAN termination applications.

Each configured FPE based PW-Port is associated with two internal SDPs (one in each direction) whose id(s) are allocated from the configured sdp-id-range.

When the FPE is associated to VXLAN termination, an internal SDP is allocated from the configured sdp-id-range and is used for R-VPLS services that terminate VXLAN IPv6. A spoke-sdp per VXLAN IPv6 R-VPLS service is created on that SDP for egress processing of the packets. Sdp-id-range cannot be modified if any of its IDs are currently in use.

Default no sdp-id-range

Parameters **from** *id* — Specifies the start of the SDP ID range (inclusive).

Values 1 to 17407

to *id* — Specifies the end of the SDP ID range.

Values 1 to 17407

2.20.2.13 APS Commands

aps

Syntax	aps
Context	config>port
Description	This command configures APS (Automatic Protection Switching). APS is used by SONET/SDH add/drop multiplexers (ADMs) or other SONET/SDH-capable equipment to protect against circuit or equipment failure. An APS group contains a working and a protect circuit and can span a single node (SC-APS) or two nodes (MC-APS). The working and protection configurations on the 7750 SRs must match the circuit configurations on the peer. This means that the working circuit on the 7750 SR must be connected to the peer's working circuit and the protect circuit must be connected to the peer's protection circuit. The aps command is only available for APS groups and not physical ports.

advertise-interval

Syntax	advertise-interval <i>advertise-interval</i> no advertise-interval
Context	config>port>aps
Description	This command specifies the time interval, in 100s of milliseconds, between 'I am operational' messages sent by both protect and working circuits to their neighbor for multi-chassis APS. The advertise-interval value is valid only for a multi-chassis APS as indicated by the value of the neighbor command value if it is not set to 0.0.0.0.
Default	10
Parameters	<i>advertise-interval</i> — Specifies the time interval, in 100s of milliseconds, between 'I am operational' messages sent by both protect and working circuits to their neighbor for multi-chassis APS. Values 10 to 650

hold-time

Syntax	hold-time <i>hold-time</i>

no hold-time

Context	config>port>aps
Description	This command specifies how much time can pass, in 100s of milliseconds, without receiving an advertise packet from the neighbor before the multi-chassis signaling link is considered not operational. The hold-time is usually 3 times the value of the advertise-interval. The value of the advertise-interval is valid only for a multi-chassis APS as indicated by the value of neighbor IP address if it is not set to 0.0.0.0.
Parameters	<i>hold-time</i> — Specifies how long to wait for an APS advertisement packet before the peer in a Multi-Chassis APS group is considered operationally down. Values 10 to 650

hold-time-aps

Syntax	hold-time-aps [Isignal-failure <i>sf-time</i>] [Isignal-degrade <i>sd-time</i>] no hold-time-aps
Context	config>port>aps
Description	This command configures hold-down timers to debounce signal failure conditions (lais, b2err-sf) and signal degrade conditions (b2err-sd) for Uni 1+1 Sig+Data APS switching mode (switching mode uni-1plus1). The no version of this command resets the hold-down timer to the default value.
Default	0 (disabled)
Parameters	<i>sf-time</i> — Specifies an integer to define the signal failure hold-down time in milliseconds. Values 1 to 100 <i>sd-time</i> — Specifies an integer to define the signal degrade hold-down time in milliseconds. Values 1 to 100

mode-annexb

Syntax	[no] mode-annexb
Context	config>port>aps
Description	This command configures the APS group for 1+1 Optimized operation as described in Annex B of ITU.T G.841. Note that Annex B operates in non-revertive bi-directional switching mode only as defined in G.841.

neighbor

Syntax	neighbor <i>ip-address</i> no neighbor
Context	config>port>aps
Description	This command specifies the neighbor's IP address only on a multi-chassis APS where the working and protect circuits are configured on different routers. When the value the neighbor IP address is set to 0.0.0.0, this implies that the APS group is configured as a single-chassis APS group. The route to the neighbor must not traverse the multi-chassis APS member (working or protect) circuits. It is recommended that the neighbor IP address configured is on a shared network between the routers that own the working and protect circuits. By default no neighbor address is configured and both the working and protect circuits should be configured on the same router (i.e., single-chassis APS). APS is assumed to be configured wholly on a single chassis.
Parameters	<i>ip-address</i> — Specifies the neighbor's IP address only on a multi-chassis APS where the working and protect circuits are configured on different routers. The node should be connected with a direct interface to ensure optimum fail-over time. Values ipv4-address: a.b.c.d ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x:-[0 to FFFF]H d: [0 to 255]D

protect-circuit

Syntax	protect-circuit <i>port-id</i> no protect-circuit
Context	config>port>aps
Description	This command configures a physical port that will act as the protection circuit for this APS group. The protect circuit port must contain only the default configuration and cannot belong to another APS group. The protect circuit port must be of the same type as the working circuit for the APS group, for the port to be added to an APS group port. If that's not the case, the command will return an error. A protection circuit can only be added if the working circuit already exists; the protection circuit must be removed from the configuration before the working circuit is removed.

When a port is a protect-circuit of an APS group, the configuration options available in the **config>port *port-id*>sonet-sdh** context is not allowed for that port unless it is part of the noted exceptions. The exception list includes these SONET/SDH commands:

- clock-source
- [no] loopback
- [no] report-alarm
- section-trace
- [no] threshold

When is port configured as a protection circuit of an APS group, the configurations described above and all service configurations related to APS port are operationally inherited by the protect circuit. If the protect circuit cannot inherit the configurations (due to resource limitations), the configuration attempt fails and an error is returned to the user.

The protect circuit must be shutdown before it can be removed from the APS group port. The inherited configuration for the circuit and APS operational commands for that circuit are not preserved when the circuit is removed from the APS group.

The **no** form of this command removes the protect-circuit.

Parameters *port-id* — Specifies the physical port that will act as the protection circuit for this APS group in the following format.

<i>port-id</i>	<i>slot/mda/port</i>		
	eth-sat-id	<i>esat-id/slot/port</i>	
		esat	keyword
		<i>id</i>	1 to 20
	pxc-id	<i>pxc-id.sub-port</i>	
		pxc	keyword
		<i>id</i>	1 to 64
		<i>sub-port</i>	a, b

Also see [Modifying Hold-Down Timer Values](#) for information about modifying the timer defaults in the event of communication delays between the APS controllers.

rdi-alarms

Syntax	rdi-alarms [suppress circuit]
Context	config>port>aps
Description	This command configures how RDI alarms (line, path, section) are generated on physical circuits of an APS ports. The command configuration changes are supported only for switching-mode set to uni_1plus1. The configuration can be changed only when no working and protecting circuit has been added. Options:

- circuit-RDI alarms are H/W-generated independently on each working and protect circuit based on RX failure of that circuit regardless of APS line status.
- suppress-RDI H/W generation on working and protect circuits is suppressed. No alarms are generated on RX failure of that circuit.

Default rdi-alarms circuit

revert-time

Syntax **revert-time** *minutes*
no revert-time

Context config>port>aps

Description This command configures the revert-time timer to determine how long to wait before switching back to the working circuit after that circuit has been restored into service.

A change in the *minutes* value takes effect upon the next initiation of the wait to restore (WTR) timer. It does not modify the length of a WTR timer that has already been started. The WTR timer of a non-revertive switch can be assumed to be infinite.

The **no** form of this command restores the default (non-revertive mode).

Default The default is to not revert back unless the protect circuit fails or there is an operator intervention.

Parameters *minutes* — Specifies the time, in minutes, to wait before reverting back to the original working circuit after it has been restored into service.

Values 0 to 60 minutes

Default 5

switching-mode

Syntax **switching-mode** {**bi-directional** | **uni-directional** | **uni-1plus1**}

Context config>port>aps

Description This command configures the switching mode for the APS group.

Parameters **bi-directional** — Configures the group to operate in Bidirectional 1+1 Signaling APS mode.

uni-directional — Configures the group to operate in Unidirectional 1+1 Signaling APS mode.

uni-1plus1 — Configures the group to operate in Unidirectional 1+1 Signaling and Datapath APS mode (7750 SR-c4/c12 platforms only).

working-circuit

Syntax	working-circuit <i>port-id</i> [number <i>number</i>] no work-circuit [number <i>number</i>]																																		
Context	config>port>aps																																		
Description	This command configures a physical port that will act as the working circuit for this APS group. The working circuit port must contain only the default configuration and cannot be part of another APS group. The working circuit must be created before the protection circuit. When a port is a working circuit of an APS group, the configuration available under config>port <i>port-id</i> context (including submenus) is not allowed for that port unless it is a part of the noted exceptions. When a port is being configured as a working circuit of an APS group, all common configuration as described above and all service configurations related to the APS port is operationally inherited by the working circuit from the <i>aps-group-id</i>. If the working circuit cannot inherit that configuration, for example, due to resource limitations, the configuration attempt fails and an error is returned to the user. Before a working circuit can be removed from an APS group, the working circuit port must be shutdown. The inherited configuration for the circuit and APS operational commands for that circuit are not preserved when the circuit is removed from the APS group. Note that all configurations for <i>aps-group-id</i> under the config>port context and its submenus and all configuration for services that use this <i>aps-group-id</i> is preserved as a non-activated configuration since the APS group no longer has any physical circuits assigned. The no form of this command removes the working-circuit. The working circuit can only be removed from the configuration after the protect circuit has been removed.																																		
Parameters	<i>port-id</i> — Specifies the physical port that will act as the working circuit for this APS group in the following format: <table><tr><td><i>port-id</i></td><td><i>slot/mda/port</i></td><td></td><td></td></tr><tr><td></td><td>eth-sat-id</td><td>esat-id/slot/port</td><td></td></tr><tr><td></td><td></td><td>esat</td><td>keyword</td></tr><tr><td></td><td></td><td>id</td><td>1 to 20</td></tr><tr><td></td><td>pxc-id</td><td>pxc-id.sub-port</td><td></td></tr><tr><td></td><td></td><td>pxc</td><td>keyword</td></tr><tr><td></td><td></td><td>id</td><td>1 to 64</td></tr><tr><td></td><td></td><td>sub-port</td><td>a, b</td></tr></table>			<i>port-id</i>	<i>slot/mda/port</i>				eth-sat-id	esat-id/slot/port				esat	keyword			id	1 to 20		pxc-id	pxc-id.sub-port				pxc	keyword			id	1 to 64			sub-port	a, b
<i>port-id</i>	<i>slot/mda/port</i>																																		
	eth-sat-id	esat-id/slot/port																																	
		esat	keyword																																
		id	1 to 20																																
	pxc-id	pxc-id.sub-port																																	
		pxc	keyword																																
		id	1 to 64																																
		sub-port	a, b																																
	<i>number</i> — Specifies the APS channel number; value is 1 or 2.																																		

Modifying Hold-Down Timer Values

Note that for APS configurations, the **hold-time down** and **hold-time up** default values are 100 ms and 500 ms respectively. But, if there is a large difference in the transmission delay between the APS working (**working-circuit**) and protect line (**protect-circuit**), it is highly recommended that you increase the default timer on the working line accordingly with the transmission delay present on the protect line.

The following output shows an example of the timers on POS interfaces.

```
A:NS044050253# show port aps-1
=====
SONET/SDH Interface
=====
Description          : APS Group
Interface            : aps-1                Speed          : oc3
Admin Status         : up                   Oper Status     : up
Physical Link        : Yes                  Loopback Mode   : none
Single Fiber Mode    : No
Clock Source         : node                 Framing         : sonet
Last State Change    : 04/11/2007 13:53:01 Port IfIndex    : 1358987264
J0 String            : 2/1/5 7750-SR-7       Section Trace Mode : string
Rx S1 Byte           : 0x00 (stu)           Rx K1/K2 Byte   : 0x00/0x00
Tx S1 Byte           : 0x0f (dnu)           Tx DUS/DNU      : disabled
Rx J0 String (Hex)   : 81 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Cfg Alarm            : loc lais lrdi ss1f lb2er-sd lb2er-sf slof slos lrei
Alarm Status         :
Hold time up         : 500 milliseconds
Hold time down       : 100 milliseconds
=====
Port Statistics
=====
Input                Output
-----
Packets              6670498              3804661
Discards              0
Unknown Proto Discards 0
=====
A:NS044050253#
```

For unprotected port these timer are different:

```
A:NS044050253# show port 2/2/2
=====
SONET/SDH Interface
=====
Description          : OC-48 SONET/SDH
Interface            : 2/2/2                Speed          : oc48
Admin Status         : up                   Oper Status     : up
Physical Link        : Yes                  Loopback Mode   : none
Single Fiber Mode    : No
APS Group            : none                 APS Role        : none
Clock Source         : loop                 Framing         : sonet
Last State Change    : 04/11/2007 14:53:53 Port IfIndex    : 37814272
J0 String            : 0x01                 Section Trace Mode : byte
Rx S1 Byte           : 0x00 (stu)           Rx K1/K2 Byte   : 0x00/0x00
Tx S1 Byte           : 0x0f (dnu)           Tx DUS/DNU      : disabled
Rx J0 String (Hex)   : af 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Cfg Alarm            : loc lrdi lb2er-sf slof slos
Alarm Status         :
```

```

Hold time up      : 500 milliseconds
Hold time down    : 0 milliseconds
Transceiver Data

Transceiver Type   : SFP
Model Number       : SFP-OC48-SR1
Transceiver Code   : OC48 SR
Laser Wavelength   : 1310
Connector Code     : LC
Manufacture date    : 2004/08/20 00:00:00
Serial Number      : 6331000705
Part Number        : CT2-MS1LBD32Z2
Optical Compliance*: 00:01:00:00:00:00:00:00
Link Len 9u        : 2 kms
Link Len 9u        : 20 * 100m
Link Len 50u       : 0 * 10m
Diag Capable       : yes
Vendor OUI         : 00:01:9c
Media              : SONET/SDH
Link Len Cu        : 0 m
Link Len 62.5u     : 0 * 10m

=====
Port Statistics
=====
                        Input                Output
-----
Packets                3870094                6656408
Discards                 0                      0
Unknown Proto Discards   0
=====
A:NS044050253#

```

wtr-annexb

Syntax	wtr-annexb <i>minutes</i>
Context	config>port>aps
Description	This command waits to restore for Annex B mode operation. The delay after which the newly active section becomes the primary section after a switch-over from the primary section to the secondary section occurs and the switch request clears normally.
Parameters	<i>minutes</i> — Specifies the time, in minutes, to wait to restore for Annex B mode operation.
Values	0 to 60
Default	5

2.20.2.14 Ethernet Port Commands

ethernet

Syntax	ethernet
Context	config>port

Description This command enables access to configure Ethernet port attributes.

This context can only be used when configuring Fast Ethernet, gigabit, or 10-Gb Ethernet LAN ports on an appropriate MDA.

mode

Syntax **mode {access | network | hybrid}**
no mode

Context config>port>ethernet
config>port>sonet-sdh>path
config>port>tdm>ds1>channel-group
config>port>tdm>ds3
config>port>tdm>e1>channel-group
config>port>tdm>e3

Description This command configures an Ethernet port for access, network, or hybrid mode of operation. It also configures a TDM channel or SONET/SDH path (sub-port) for access or network mode operation.

An access port or channel is used for customer facing traffic on which services are configured. A Service Access Point (SAP) can only be configured on an access port or channel. When a port is configured for access mode, the appropriate encap-type must be specified to distinguish the services on the port or SONET path. Once an Ethernet port, a TDM channel or a SONET path has been configured for access mode, multiple services can be configured on the Ethernet port, a TDM channel or SONET path. Note that ATM, Frame Relay, and cHDLC port parameters can only be configured in the access mode.

A network port or channel participates in the service provider transport or infrastructure network when a network mode is selected. When the network option is configured, the encap-type cannot be configured for the port/channel.

When network mode is selected on a SONET/SDH path, the appropriate control protocols are activated when the need arises. For example, configuring an IP interface on the SONET path activates IPCP while the removal of the IP interface causes the IPCP to be removed. The same applies for MPLS, MPLSCP, and OSICP. When configuring a SONET/SDH port, the mode command must be entered in the channel context or an error message is generated.

A hybrid Ethernet port allows the combination of network and access modes of operation on a per-VLAN basis and must be configured as either dot1q or QinQ encapsulation.

When the hybrid port is configured to the dot1q encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and an unused VLAN tag value. The format is <port-id>:qtag1. A SAP of format <port-id>:* also supported.

The user configures a network IP interface under `config>router>if>port` by providing the port name which consists of the port-id of the hybrid mode port and an unused VLAN tag value. The format is `<port-id>:qtag1`. The user must explicitly enter a valid value for qtag1. The `<port-id>.*` value is not supported on a network IP interface. The 4096 VLAN tag space on the port is shared among VLAN SAPs and VLAN network IP interfaces.

When the hybrid port is configured to QinQ encapsulation, the user configures a SAP inside a service simply by providing the SAP ID which must include the port-id value of the hybrid mode port and the outer and inner VLAN tag values. The format is `<port-id>:qtag1.qtag2`. A SAP of format `<port-id>:qtag1.*` is also supported. The outer VLAN tag value must not have been used to create an IP network interface on this port. In addition, the qtag1.qtag2 value combination must not have been used by another SAP on this port.

The user configures a network IP interface under `config>router>if>port` by providing the port name which consists of the port-id of the hybrid mode port and a VLAN tag value. The format is `<port-id>:qtag1.*`. An outer VLAN tag qtag2 of * creates an IP network interface. In addition, the qtag1.qtag2 value combination must not have been used on another SAP or IP network interface on this port.

The **no** form of this command restores the default.

Default mode network — For Ethernet ports.

mode access — For TDM channel or SONET paths.

Parameters **access** — Configures the Ethernet port, TDM channel or SONET path as service access.

network — Configures the Ethernet port, TDM channel or SONET path for transport network use.

hybrid — Configures the Ethernet port for hybrid use.

access

Syntax **access**

Context `config>port>ethernet`

Description This command configures Ethernet access port parameters.

egress

Syntax **egress**

Context `config>port>ethernet>access`
`config>port>ethernet>network`

Description This command configures Ethernet access egress port parameters.

queue-group

Syntax	queue-group <i>queue-group-name</i> [instance <i>instance-id</i>] [create] no queue-group <i>queue-group-name</i> [instance <i>instance-id</i>]
Context	config>port>ethernet>access>egress config>port>ethernet>access>ingress
Description	This command creates an ingress or egress queue group on an Ethernet port. A queue group is a collection of queues identified by a group name. Queue groups created on access ports are used as an alternative queue destination for SAPs. Within a SAP, a forwarding class may be redirected from the local SAP queue to a port queue group queue. The forwarding classes from multiple SAPs may be redirected to the same queue group which can be used to minimize the number of per-SAP queues. Queue groups may be created on both access and network oriented ports. When the port is in access mode, the queue groups must be created within the port access node. Within the access node, queue groups are also configured as ingress or egress. Access ingress queue groups can only be used by ingress SAP forwarding classes and only a single ingress queue group per port is supported. Multiple access egress queue groups may be created on a single port and are used by egress SAP forwarding classes. The instance-id parameter identifies different instances of the same queue group template. Creating multiple queue groups with a different instance ID but the same queue group name results in separate queue groups being created on the port. The instance-id parameter is only valid for egress queue groups on access ports. When the queue group is created in an ingress port context, the group-name must be an existing ingress queue group template. Similarly, queue groups created in an egress port context must have a group-name of an existing egress queue group template. Two ingress queue groups with the same name cannot be created on the same port. Two egress queue groups can only be created on the same port with the same queue group template name if they have different instance-id values. The queues defined in the template are created on the queue group. The queue parameters within the template are used as the default queue parameters for each queue in the queue group. The default queue parameters for each queue may be overridden on the queue group with specific queue parameters. Each queue group supports the application of a scheduler-policy for the purpose of managing the queues within the group into an aggregate SLA. The queues defined within the template may be configured with parent scheduler defining the mapping of a queue to one of the schedulers within the scheduler policy. Egress queue groups also support the agg-rate parameter and the queues in the egress template support the port-parent command. Each command is used for configuring egress port virtual scheduling behavior. Each queue group allows the application of an accounting policy and the ability to enable and disable collecting statistics. The statistics are derived from the queue counters on each queue within the queue group. The accounting policy defines which queue counters are collected and to which accounting file they will be written.

A queue group does not have an administrative shutdown or no shutdown command. A queue group is considered to be always on once created.

When creating a queue group, the system will attempt to allocate queue resources based on the queues defined in the queue group template. If the appropriate queue resources do not currently exist, the queue group will not be created. Ingress port queue groups do not support the shared-queuing or multipoint-shared queuing behavior.

When the queue group is created on a LAG (Link Aggregation Group), it must be created on the primary port member. The primary port member is the port with the lowest port ID based on the slot, MDA position and port number on the MDA. A queue group created on the primary LAG port will be automatically created on all other port members. If a new port is being added to a LAG with an existing queue group, the queue group must first be created on the port prior to adding the port to the LAG. If the LAG queue group has queue overrides, the queue overrides must also be defined on the port queue group prior to adding the port to the LAG.

A port queue group cannot be removed from the port when a forwarding class is currently redirected to the group. All forwarding class redirections must first be removed prior to removing the queue group.

- Parameters**
- group-name* — The group-name parameter is required when executing the port queue-group command. The specified group-name must exist as an ingress or egress queue group template depending on the ingress or egress context of the port queue group. Only a single queue group may be created on an ingress port. Multiple queue groups may be created on an egress port.
 - instance-id* — Specifies the identification of a specific instance of the egress queue-group. This parameter is only valid for egress access port queue groups.
- Values** 1 to 65535
- create** — Keyword used to associate the queue group. The create keyword requirement can be enabled/ disabled in the environment>create context.

egress

- Syntax** **egress**
- Context** config>port>ethernet
- Description** This command configures Ethernet egress port parameters.

ingress

- Syntax** **ingress**
- Context** config>port>ethernet>access
- Description** This command configures Ethernet access ingress port parameters.

queue-group

Syntax	[no] queue-group <i>queue-group-name</i> [instance <i>instance-id</i>] [create]
Context	config>port>ethernet>access>egr config>port>ethernet>access>ing
Description	This command creates an ingress or egress queue group on an Ethernet port. A queue group is a collection of queues identified by a group name. Queue groups created on access ports are used as an alternative queue destination for SAPs. Within a SAP, a forwarding class may be redirected from the local SAP queue to a port queue group queue. The forwarding classes from multiple SAPs may be redirected to the same queue group which can be used to minimize the number of per-SAP queues. Queue groups may be created on both access and network oriented ports. When the port is in access mode, the queue groups must be created within the port access node. Within the access node, queue groups are also configured as ingress or egress. Access ingress queue groups can only be used by ingress SAP forwarding classes and only a single ingress queue group per port is supported. Multiple access egress queue groups may be created on a single port and are used by egress SAP forwarding classes. The instance-id parameter identifies different instances of the same queue group template. Creating multiple queue groups with a different instance ID but the same queue group name results in separate queue groups being created on the port. The instance-id parameter is only valid for egress queue groups on access ports. When the queue group is created in an ingress port context, the group-name must be an existing ingress queue group template. Similarly, queue groups created in an egress port context must have a group-name of an existing egress queue group template. Two ingress queue groups with the same name cannot be created on the same port. Two egress queue groups can only be created on the same port with the same queue group template name if they have different instance-id values. The queues defined in the template are created on the queue group. The queue parameters within the template are used as the default queue parameters for each queue in the queue group. The default queue parameters for each queue may be overridden on the queue group with specific queue parameters. Each queue group supports the application of a scheduler-policy for the purpose of managing the queues within the group into an aggregate SLA. The queues defined within the template may be configured with parent scheduler defining the mapping of a queue to one of the schedulers within the scheduler policy. Egress queue groups also support the agg-rate parameter and the queues in the egress template support the port-parent command. Each command is used for configuring egress port virtual scheduling behavior. Each queue group allows the application of an accounting policy and the ability to enable and disable collecting statistics. The statistics are derived from the queue counters on each queue within the queue group. The accounting policy defines which queue counters are collected and to which accounting file they will be written.

A queue group does not have an administrative shutdown or no shutdown command. A queue group is considered to be always on once created.

When creating a queue group, the system will attempt to allocate queue resources based on the queues defined in the queue group template. If the appropriate queue resources do not currently exist, the queue group will not be created. Ingress port queue groups do not support the shared-queuing or multipoint-shared queuing behavior.

When the queue group is created on a LAG (Link Aggregation Group), it must be created on the primary port member. The primary port member is the port with the lowest port ID based on the slot, MDA position and port number on the MDA. A queue group created on the primary LAG port will be automatically created on all other port members. If a new port is being added to a LAG with an existing queue group, the queue group must first be created on the port prior to adding the port to the LAG. If the LAG queue group has queue overrides, the queue overrides must also be defined on the port queue group prior to adding the port to the LAG.

A port queue group cannot be removed from the port when a forwarding class is currently redirected to the group. All forwarding class redirections must first be removed prior to removing the queue group.

- Parameters**
- queue-group-name* — The group-name parameter is required when executing the port queue-group command. The specified group-name must exist as an ingress or egress queue group template depending on the ingress or egress context of the port queue group. Only a single queue group may be created on an ingress port. Multiple queue groups may be created on an egress port.
 - instance-id* — Specifies the identification of a specific instance of the queue-group.
- Values** 1 to 65535
- create** — Keyword used to associate the queue group. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

agg-rate

- Syntax** [no] **agg-rate**
- Context** config>port>ethernet>access>egr>qgrp
config>port>ethernet>access>egr>vport
config>port>ethernet>network>egr>qgrp
- Description** This command controls an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, and **queue-frame-based-accounting**.
- When specified under a VPORT, the agg-rate rate, port-scheduler-policy and scheduler-policy commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command.

limit-unused-bandwidth

Syntax	[no] limit-unused-bandwidth
Context	<pre>config>port>ethernet>access>egr>qgrp>agg-rate config>port>ethernet>access>egr>vport>agg-rate config>port>ethernet>network>egr>qgrp>agg-rate config>port>sonet-sdh>path>access>egress>vport</pre>
Description	This command enables aggregate rate overrun protection on the agg-rate context. The no form of the command disables aggregate rate overrun protection on the agg-rate context.

queue-frame-based-accounting

Syntax	[no] queue-frame-based-accounting
Context	<pre>config>port>ethernet>access>egr>qgrp>agg-rate config>port>ethernet>access>egr>vport>agg-rate config>port>ethernet>network>egr>qgrp>agg-rate config>port>sonet-sdh>path>access>egress>vport</pre>
Description	This command enables frame based accounting on all policers and queues associated with the agg-rate context. It is only supported on Ethernet ports but not on HSMDA Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame-based accounting is configured, regardless of how offsets are applied to the statistics. The no form of the command disables frame based accounting on all policers and queues associated with the agg-rate context.

rate

Syntax	rate <i>kilobits-per-second</i> no rate
Context	<pre>config>port>ethernet>access>egr>qgrp>agg-rate config>port>ethernet>access>egr>vport>agg-rate config>port>ethernet>network>egr>qgrp>agg-rate config>port>sonet-sdh>path>access>egress>vport</pre>
Description	This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, VPORT and so on).
Parameters	<i>kilobits-per-second</i> — Specifies the rate limit for the VPORT, in kb/s. Values 1 to 3200000000 or max

host-match

Syntax	host-match dest <i>destination-string</i> [create] no host-match dest <i>destination-string</i>
Context	config>port>ethernet>access>egr>qgrp
Description	This command configures host matching for the Ethernet port egress queue-group. The no form of the command removes host matching for the Ethernet port egress queue-group.
Parameters	<i>destination-string</i> — Specifies a host match destination string up to 32 characters. create — Keyword used to create the host match. The create keyword requirement can be enabled or disabled in the environment>create context.

hsmda-queue-override

Syntax	[no] hsmda-queue-override
Context	config>port>ethernet>access>egress>queue-group
Description	This command configures HSMDA egress queue overrides parameters.

packet-byte-offset

Syntax	packet-byte-offset { add <i>add-bytes</i> subtract <i>sub-bytes</i> } no packet-byte-offset
Context	config>port>ethernet>access>egress>queue-group>hsmda-queue-override
Description	This command configures a packet offset for HSMDA queue accounting.
Parameters	<i>add-bytes</i> — Specifies the number of bytes to add. Values 0 to 31 <i>sub-bytes</i> — Specifies the number of bytes to subtract. Values 1 to 64

queue

Syntax	queue <i>queue-id</i> [create] queue <i>queue-id</i>
Context	config>port>ethernet>access>egress>queue-group>hsmda-queue-override

Description This command configures overrides for an HSMDA egress queue.

Parameters *queue-id* — Specifies the queue ID.

Values 1 to 8

create — Keyword used to create

mbs

Syntax **mbs** {[0 to 2625] kilobytes | [0 to 268800] bytes | default}
no mbs

Context config>port>ethernet>access>egress>queue-group>hsmda-queue-override>queue

Description This command configures the maximum buffer space override.

rate

Syntax **rate** *pir-rate*
no rate

Context config>port>ethernet>access>egress>queue-group>hsmda-queue-override>queue

Description This command configures the peak information rate.

Parameters *pir-rate* — Specifies the pir weight.

Values 1 to 100000000, or max

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>port>ethernet>access>egress>queue-group>hsmda-queue-override>queue

Description This command configures the slope policy.

Parameters *hsmda-slope-policy-name* — Specifies the HSMDA slope policy name, up to 32 characters.

wrr-weight

Syntax **wrr-weight** *weight*
no wrr-weight

Context	config>port>ethernet>access>egress>queue-group>hsmda-queue-override>queue
Description	This command configures the weighted round robin (WRR).
Parameters	<i>weight</i> — Specifies the WRR weight value.
Values	1 to 32

secondary-shaper

Syntax	secondary-shaper <i>secondary-shaper-name</i> no secondary-shaper
Context	config>port>ethernet>access>egress>queue-group>hsmda-queue-override
Description	This command configures the HSMDA egress secondary shaper.
Parameters	<i>secondary-shaper-name</i> — Specifies the HSMDA secondary shaper name, up to 32 characters.

wrr-policy

Syntax	wrr-policy <i>wrr-policy-name</i>
Context	config>port>ethernet>access>egress>queue-group>hsmda-queue-override
Description	This command configures the HSMDA egress wrr-policy.
Parameters	<i>wrr-policy-name</i> — Specifies the HSMDA egress WRR policy name, up to 32 characters.

policer-control-policy

Syntax	policer-control-policy <i>policy-name</i> no policer-control-policy
Context	config>port>ethernet>network>egress>queue-group>policer-control-policy
Description	This command configures the policer control policy for the QoS egress queue-group.
Parameters	<i>policy-name</i> — Specifies the name of the policer control policy, up to 32 characters.

queue-overrides

Syntax	queue-overrides
Context	config>port>ethernet>access>egr>qgrp


```
config>port>ethernet>access>ing>qgrp
config>port>ethernet>network>egr>qgrp
```

Description This command enables the context to define optional queue parameter overrides for each queue within the queue group.

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>port>ethernet>access>egr>qgrp>qover
config>port>eth>network>egr>qgrp>qover

Description This command associates a queue for use in a queue group template. The defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object which is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template's name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP egress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

The **no** form of the command removes the queue-id from the configuration.

Parameters *queue-id* — Specifies the queue ID.

Values 1 to 8

create — Mandatory when creating an entry.

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>port>ethernet>access>ing>qgrp>qover

Description This command associates a queue for use in a queue group template. The defined queue-id acts as a repository for the default parameters for the queue. The template queue is created on each queue-group object which is created with the queue group template name. Each queue is identified within the template by a queue-id number. The template ensures that all queue groups created with the template's name will have the same queue-ids providing a uniform structure for the forwarding class redirection commands in the SAP egress QoS policies. The parameters within the template queue will be used as the default settings for each queue in the actual queue group. The queue parameters may be individually changed for each queue in each queue group using per queue overrides.

The **no** form of the command removes the queue-id from the configuration.

- Parameters** *queue-id* — Specifies the queue ID.
- Values** 1 to 32
- create** — Mandatory when creating an entry.

parent

- Syntax** **parent** [**weight** *weight*] [**cir-weight** *cir-weight*]
no parent
- Context** config>port>ethernet>access>egr>qgrp>qover>q
- Description** This command, when used in the *queue-overrides* context for a queue group queue, defines an optional **weight** and **cir-weight** for the queue treatment by the parent scheduler that further governs the available bandwidth given the queue aside from the queue PIR setting. When multiple schedulers and/or queues share a child status with the parent scheduler, the weight or level parameters define how this queue contends with the other children for the parent bandwidth.
- Parameters** *weight* — Weight defines the relative weight of this queue in comparison to other child schedulers and queues while vying for bandwidth on the parent scheduler-name. Any queues or schedulers defined as weighted receive no parental bandwidth until all strict queues and schedulers on the parent have reached their maximum bandwidth or are idle. In this manner, weighted children are considered to be the lowest priority.
- Values** 0 to 100
- Default** 1
- cir-weight* — Defines the weight the queue will use at the within-cir port priority level. The weight is specified as an integer value from 0 to 100 with 100 being the highest weight. When the cir-weight parameter is set to a value of 0 (the default value), the queue or scheduler does not receive bandwidth during the port schedulers within-cir pass and the cir-level parameter is ignored. If the cir-weight parameter is 1 or greater, the cir-level parameter comes into play.
- Values** 0 to 100

adaptation-rule

- Syntax** **adaptation-rule** [**pir** *adaptation-rule*] [**cir** {**max** | **min** | **closest**}]
no adaptation-rule
- Context** config>port>ethernet>access>egr>qgrp>qover>q
config>port>ethernet>access>ing>qgrp>qover>q
config>port>ethernet>network>egr>qover>q

Description	This command specifies the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint. This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case the configuration of the adaptation rule is performed under the hs-wrr-group within the egress queue group template. The no form of the command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific adaptation-rule is removed, the default constraints for rate and cir apply.
Default	adaptation-rule pir closest cir closest
Parameters	pir — Defines the constraints enforced when adapting the PIR rate defined within the queue queue-id rate command. The pir parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the rate command is not specified, the default applies. cir — Defines the constraints enforced when adapting the CIR rate defined within the queue queue-id rate command. The cir parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the cir parameter is not specified, the default constraint applies. adaptation-rule — Specifies the adaptation rule to be used while computing the operational CIR or PIR value. Values max — The max (maximum) option is mutually exclusive with the min and closest options. When max is defined, the operational PIR for the queue will be equal to or less than the administrative rate specified using the rate command. min — The min (minimum) option is mutually exclusive with the max and closest options. When min is defined, the operational PIR for the queue will be equal to or greater than the administrative rate specified using the rate command. closest — The closest parameter is mutually exclusive with the min and max parameter. When closest is defined, the operational PIR for the queue will be the rate closest to the rate specified using the rate command.

burst-limit

Syntax	burst-limit {default size [bytes kilobytes]} no burst-limit
Context	config>port>ethernet>access>egr>qgrp>qover>q

Description	The queue burst-limit command overrides the shaping burst size for a queue. The configured size defines the shaping leaky bucket threshold level that indicates the maximum burst over the queue's shaping rate. The no form of this command removes the current burst limit override for the queue. The queue's burst limit is controlled by its defining template.
Default	no burst-limit
Parameters	default — Reverts the queue's burst limit to the system default value. size — When a numeric value is specified (<i>size</i>), the system interprets the value as an explicit burst limit size. The value is expressed as an integer and, by default, is interpreted as the burst limit in kilobytes. If the value is intended to be interpreted in bytes, the bytes qualifier must be added following <i>size</i>. Values 1 to 13671 kilobytes 1 to 14000000 bytes Default No default for <i>size</i>; use the default keyword to specify default burst limit. bytes — Specifies that the value given for <i>size</i> must be interpreted as the burst limit in bytes. kilobytes — Specifies that the value given for <i>size</i> must be interpreted as the burst limit in kilobytes. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

cbs

Syntax	cbs <i>size-in-kbytes</i> no cbs
Context	config>port>ethernet>access>egr>qgrp>qover>q config>port>ethernet>access>ing>qgrp>qover>q config>port>ethernet>network>egr>qover>q
Description	This command defines the default committed buffer size for the template queue. Overall, the cbs command follows the same behavior and provisioning characteristics as the cbs command in the queue-group or network QoS policy. The exception is the addition of the cbs-value qualifier keywords bytes or kilobytes. The no form of this command restores the default CBS size to the template queue.
Default	cbs default

Parameters	<i>size-in-kbytes</i> — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).
Values	0 to 1048576 or default

drop-tail

Syntax	drop-tail
Context	config>port>eth>access>ing>qgrp>qover>q config>port>eth>access>egr>qgrp>qover>q config>port>ethernet>network>egr>qgrp>qover>q
Description	This command enters the context to configure queue drop tail parameters.

low

Syntax	low
Context	config>port>eth>access>ing>qgrp>qover>q>drop-tail config>port>eth>access>egr>qgrp>qover>q>drop-tail config>port>ethernet>network>egr>qgrp>qover>q>drop-tail
Description	This command enters the context to configure the queue low drop tail parameters. The low drop tail defines the queue depth beyond which out-of-profile packets will not be accepted into the queue and will be discarded.

percent-reduction-from-mbs

Syntax	percent-reduction-from-mbs <i>percent</i> no percent-reduction-from-mbs
Context	config>port>ethernet>access>egr>qgrp>qover>q>drop-tail>low config>port>ethernet>access>ing>qgrp>qover>q>drop-tail>low config>port>ethernet>network>egr>qgrp>qover>q>drop-tail>low
Description	This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail will be at 420 kbytes and out-of-profile packets will not be accepted into the queue if its depth is greater than this value, and so will be discarded.
Parameters	<i>percent</i> — Specifies the percentage reduction from the MBS for a queue drop tail.
Values	0 to 100, default

mbs

Syntax	mbs { <i>size</i> [bytes kilobytes] default } no mbs
Context	config>port>ethernet>access>egr>qgrp>qover>q config>port>ethernet>access>ing>qgrp>qover>q config>port>ethernet>network>egr>qgrp>qover>q
Description	The Maximum Burst Size (MBS) command specifies the default maximum buffer size for the template queue. The value is given in kilobytes. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueueing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue. The queue-group or network egress QoS context for mbs provides a mechanism for overriding the default maximum size for the queue. The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer will be available when needed or that the packets RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel. This command applies to egress queue group queues as the queue-delay is only supported on egress queues. This command the queue-delay command are mutually exclusive. The no form of this command returns the MBS size assigned to the queue to the value.
Default	mbs default
Parameters	size — The <i>size</i> parameter is required when specifying mbs and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional byte and kilobyte keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes. Values 0 to 1073741824 bytes — When byte is defined, the value given for size is interpreted as the queue's MBS value given in bytes. kilobytes — When kilobytes is defined, the value is interpreted as the queue's MBS value given in kilobytes. default — Specifying the keyword default sets the MBS to its default value.

monitor-depth

Syntax	[no] monitor-depth
Context	config>port>eth>access>ing>qgrp>qover>q config>port>eth>access>egr>qgrp>qover>q config>port>ethernet>network>egr>qgrp>qover>q
Description	This command enables queue depth monitoring for the specified queue. This command and the dynamic-mbs command are mutually exclusive on the related queue group queue. The no form of the command removes queue depth monitoring for the specified queue.

percent-rate

Syntax	percent-rate <i>pir-percent</i> [cir <i>cir-percent</i>]
Context	config>port>eth>access>egr>qgrp>qover>q config>port>ethernet>network>egr>qgrp>qover>q
Description	This command specifies percent rates (CIR and PIR). This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the percent-rate is performed under the hs-wrr-group within the egress queue group template.
Parameters	<i>pir-percent</i> — Specifies the PIR as a percentage. Values 0.01 to 100.00 <i>cir-percent</i> — Specifies the CIR as a percentage. Values 0.00 to 100.00

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>port>ethernet>access>egr>qgrp>qover>q config>port>ethernet>access>ing>qgrp>qover>q config>port>ethernet>network>egr>qover>q
Description	This command specifies the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

This command is ignored for egress HSQ queue group queues which are attached to an HS WRR group within an associated HS attachment policy. In this case, the configuration of the **rate** is performed under the **hs-wrr-group** within the egress queue group template.

The **no** form of the command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default	rate max cir 0 - The max default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The max value is mutually exclusive to the pir-rate value.
Parameters	<i>pir-rate</i> — Defines the administrative PIR rate, in kilobits per second, for the queue. When the rate command is executed, a valid PIR setting must be explicitly defined. When the rate command has not been executed, the default PIR of max is assumed. Fractional values are not allowed and must be given as a positive integer. The actual PIR rate is dependent on the queue's adaptation-rule parameters and the actual hardware where the queue is provisioned. Values 1 to 200000000, max Default max <i>cir-rate</i> — The cir parameter overrides the default administrative CIR used by the queue, in kilobits per second. When the rate command is executed, a CIR setting is optional. When the rate command has not been executed or the cir parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be given as a positive integer. Values 0 to 200000000, max Default 0

scheduler-override

Syntax	[no] scheduler-override
Context	config>port>ethernet>access>egr>qgrp config>port>ethernet>access>ing>qgrp

- Description** This command specifies the set of attributes whose values have been overridden by management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the ingress or egress queue group template.
- The **no** form of the command removes all of the scheduler overrides and returns the scheduler's parent weight and CIR weight, and its PIR and CIR to the values configured in the applied scheduler policy.

scheduler

- Syntax** **scheduler** *scheduler-name* [**create**]
no scheduler *scheduler-name*
- Context** config>port>ethernet>access>egr>qgrp>sched-override
config>port>ethernet>access>ing>qgrp>sched-override
- Description** This command can be used to override specific attributes of the specified scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers. The *scheduler-name* must exist in the applied scheduler policy.
- The **no** form of the command removes the scheduler overrides for the specified scheduler and returns the scheduler's parent weight and CIR weight, and its PIR and CIR to the values configured in the applied scheduler policy.
- Parameters** *scheduler-name* — Specifies the name of the scheduler.
- Values** Valid names consist of any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
- create** — Creates a new scheduler for this port.

parent

- Syntax** **parent** [**weight** *weight*] [**cir-weight** *cir-weight*]
no parent
- Context** config>port>ethernet>access>egr>qgrp>sched-override>scheduler
config>port>ethernet>access>ing>qgrp>sched-override>scheduler

Description	This command can be used to override the scheduler's parent weight and CIR weight. The weights apply to the associated level/cir-level configured in the applied scheduler policy. The scheduler name must exist in the applied scheduler policy. The override weights are ignored if the scheduler does not have a parent command configured in the scheduler policy - this allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the queue group overrides. If the parent scheduler does not exist, causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers. The no form of the command returns the scheduler's parent weight and cir-weight to the value configured in the applied scheduler policy.
Default	no parent
Parameters	<i>weight</i> — Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict level defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total distributes the available bandwidth at that level. A weight is considered to be active when the queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit. A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level. Values 0 to 100 <i>cir-weight</i> — Specifies the relative weight of this scheduler in comparison to other child schedulers and queues at the same cir-level defined by the cir-level parameter in the applied scheduler policy. Within the strict cir-level, all cir-weight values from active children at that level are summed and the ratio of each active child's cir-weight to the total distributes the available bandwidth at that level. A cir-weight is considered to be active when the policer, queue, or scheduler that the cir-weight pertains to has not reached the CIR and still has packets to transmit. A 0 (zero) cir-weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level. Values 0 to 100

rate

Syntax	rate <i>pir-rate</i> [cir <i>cir-rate</i>] no rate
Context	config>port>ethernet>access>egr>qgrp>sched-override>scheduler config>port>ethernet>access>ing>qgrp>sched-override>scheduler

Description	This command can be used to override specific attributes of the specified scheduler rate. The rate command defines the maximum bandwidth that the scheduler can offer its child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase. The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates. When a scheduler is defined without specifying a rate, the default rate is max. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child policers, queues, and schedulers to operate at their maximum rates. The no form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.
Parameters	<i>pir-rate</i> — Specifies the PIR rate. Any other value will result in an error without modifying the current PIR rate. Values 1 to 3200000000, max <i>cir-rate</i> — Specifies the CIR rate. If the CIR is set to max, then the CIR rate is set to infinity. The sum keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers, or queues. Values 0 to 3200000000, max, sum

scheduler-policy

Syntax	scheduler-policy <i>scheduler-policy-name</i> no scheduler-policy
Context	config>port>ethernet>egress>queue-group config>port>ethernet>ingress>queue-group config>port>ethernet>network>egress>queue-group
Description	This command configures a scheduler policy for the egress queue group.
Parameters	<i>scheduler-policy-name</i> — Specifies the scheduler policy name, up to 32 characters.

exp-secondary-shaper

Syntax	exp-secondary-shaper <i>secondary-shaper-name</i> [create]
Context	config>port>ethernet>egress
Description	This command configures the Ethernet egress expanded secondary shaper on this port.
Parameters	<i>secondary-shaper-name</i> — Specifies the secondary shaper name to apply to this port, up to 32 characters. create — Creates a new secondary shaper for this port.

agg-burst

Syntax	agg-burst
Context	config>port>ethernet>egress>exp-secondary-shaper
Description	This command specifies the aggregate burst limits.

high-burst-increase

Syntax	high-burst-increase <i>size</i> [bytes kilobytes] no high-burst-increase
Context	config>port>ethernet>egress>exp-secondary-shaper>agg-burst
Description	This command specifies a high burst increase.
Parameters	<i>size</i> — Specifies the shaping burst size. Values 0 to 65528 bytes — Specifies to use the size in bytes above the low burst limit to be used as the high burst threshold. kilobytes — Specifies to use the size kilobytes above the low burst limit to be used as the high burst threshold.

low-burst-limit

Syntax	low-burst-limit <i>size</i> [bytes kilobytes] no low-burst-limit
Context	config>port>ethernet>egress>exp-secondary-shaper>agg-burst
Description	This command specifies a low burst limit.

Parameters *size* — Specifies the low burst size.
 Values 1 to 327680
 bytes — Specifies to use the size in bytes.
 kilobytes — Specifies to use the size in kilobytes.

class

Syntax **class** *class-number* **rate** *rate* [**monitor-threshold** *size-in-kilobytes*] [**burst-limit** *size*] [**bytes**
 | **kilobytes**]
 no class *class-number*

Context config>port>ethernet>egress>exp-secondary-shaper

Description This command assigns the low burst maximum class to associate with the Ethernet egress expanded secondary shaper.

 The **no** form of the command returns the class id for the Ethernet egress expanded secondary shaper to the default value.

Parameters *class-number* — Specifies the class identifier of the low burst max class for the shaper.
 Values 1 to 8
 rate — Specifies the rate limit for the secondary shaper.
 Values max, 1 to 10000000 kb/s
 size-in-kilobytes — Specifies the monitor threshold for the secondary shaper.
 Values 0 to 8190
 size — Specifies the burst limit size.
 Values 1 to 327680
 bytes — Specifies to use the size in bytes.
 kilobytes — Specifies to use the size in kilobytes.

low-burst-max-class

Syntax **low-burst-max-class** *class*
 no low-burst-max-class

Context config>port>ethernet>egress>exp-secondary-shaper

Description This command specifies the class to associate with the Ethernet egress expanded secondary shaper.

The **no** form of the command returns the class number value for the Ethernet egress expanded secondary shaper to the default value.

Parameters *class* — Specifies the class number of the class for the secondary shaper.
Values 1 to 8

rate

Syntax **rate** *rate* [**monitor-threshold** *size-in-kbytes*]
no rate

Context config>port>ethernet>egress>exp-secondary-shaper

Description This command configures the shaper's metering and optional profiling rates. The metering rate is used by the system to configure the shaper's PIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's violate (PIR) threshold.

The **no** form of this command restores the default metering and profiling rate to a policer.

Parameters *rate* — Specifies the exp-secondary-shaper rate.
Values max, 1 to 10000000 kb/s
size-in-kbytes — Specifies the monitor threshold size in kbytes.
Values 0 to 8190

vport

Syntax **vport** *name* [**create**]
no vport *name*

Context config>port>ethernet>access>egress
config>port>sonet-sdh>path>access>egress

Description This command configures a scheduling node, referred to as virtual port, within the context of an egress Ethernet port. The Vport scheduler operates either like a port scheduler with the difference that multiple Vport objects can be configured on the egress context of an Ethernet port, or it can be an aggregate rate when an egress port-scheduler policy is applied to the port.

The Vport is always configured at the port level even when a port is a member of a LAG.

When a port scheduler policy is applied to a Vport the following command is used:

config>port>ethernet>access>egress>vport>port-scheduler-policy *port-scheduler-policy-name*

The CLI will not allow the user to apply a port scheduler policy to a Vport if one has been applied to the port. Conversely, the CLI will not allow the user to apply a port scheduler policy to the egress of an Ethernet port if one has been applied to any Vport defined on the access egress context of this port. The **agg-rate**, along with an egress port-scheduler, can be used to ensure that a given Vport does not oversubscribe the port's rate.

SAP and subscriber host queues can be port-parented to a Vport scheduler in a similar way they port-parent to a port scheduler or can be port-parented directly to the egress port-scheduler if the **agg-rate** is used.

Parameters *name* — Specifies the name of the Vport scheduling node and can be up to 32 ASCII characters. This does not need to be unique within the system but is unique within the port or a LAG.

agg-rate

Syntax **[no] agg-rate**

Context config>port>sonet-sdh>path>access>egress>vport
config>port>ethernet>access>egress>vport

Description This command configures an aggregate rate for the Vport. The agg-rate rate, port-scheduler-policy and scheduler-policy commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate or port-scheduler-policy involves removing the existing command and applying the new command.

egress-rate-modify

Syntax **[no] egress-rate-modify**

Context config>port>ethernet>access>egress>vport
config>port>sonet-sdh>path>access>egress>vport

Description This command applies HQoS Adjustment to a Vport. HQoS Adjustment refers to the dynamic adjustment of the rate limit at an QoS enforcement point within a Nokia router when the multicast traffic stream is disjointed from the unicast traffic stream. This QoS enforcement point within the router represents the physical point further down in the access part of the network where the two streams join each other and potentially can cause congestion.

An example would be a PON port which is shared amongst subscriber's multicast traffic (single copy of each channel) and subscriber's unicast traffic. The bandwidth control point for this PON port resides in the upstream Nokia BNG node in the form of a Vport. In the case where the multicast delivery method of the BNG utilizes redirection, the multicast traffic in the BNG will flow outside of the subscriber or the Vport context and thus will bypass any bandwidth enforcement in the Nokia router. To correct this, a Vport bandwidth adjustment is necessary in the router that will account for the multicast bandwidth consumption that is bypassing Vport in the router but is present in the PON port whose bandwidth is controlled by Vport.

An estimate of the multicast bandwidth consumption on the PON port can be made at the Vport level based on the IGMP messages sourced from the subscribers behind the PON port. This process is called HQoS Adjustment.

A multicast channel bandwidth is subtracted from or added to the Vport rate limit according to the received IGMP Join/Leave messages and the channel bandwidth definition policy associated with the Vport (indirectly through a group-interface). Since the multicast traffic on the PON port is shared amongst subscribers behind this PON port, only the first IGMP Join or the last IGMP Leave per multicast channel is tracked for the purpose of the Vport bandwidth modification.

The Vport rate that will be affected by this functionality depends on the configuration:

- In case the **agg-rate** within the Vport is configured, its value will be modified based on the IGMP activity associated with the subscriber under this Vport.
- In case the port-scheduler-policy within the Vport is referenced, the max-rate defined in the corresponding port-scheduler-policy will be modified based on the IGMP activity associated with the subscriber under this Vport.

The channel bandwidth definition policy is defined in the mcac policy in the **config>router>mcac>policy** context. The policy is applied under the group-interface or in case of redirection under the redirected-interface.

The rates in effect can be displayed with the following two commands:

show port 1/1/5 vport *name*

qos scheduler-hierarchy port *port-id vport vport-name*

The configuration of a scheduler policy under a VPORT, which is only applicable to Ethernet interfaces, is mutually exclusive with the configuration of the **egress-rate-modify** parameter.

Context: HQoS Adjustment for Vport is disabled.

host-match

Syntax	host-match dest <i>description-string</i> [create] no host-match dest <i>destination-string</i>
Context	config>port>sonet-sdh>path>access>egress>vport config>port>ethernet>access>egress>vport
Description	This command specifies the destination and organization strings to be used for matching subscriber hosts with this Vport.

The parent Vport of a subscriber host queue, which has the port-parent option enabled, is determined by matching the destination string **dest** string associated with the subscriber and the organization string **org** string associated with the subscriber host with the strings defined under a Vport on the port associated with the subscriber.

If a given subscriber host queue does not have the port-parent option enabled, it will be foster-parented to the Vport used by this subscriber and which is based on matching the dest string and org string. If the subscriber could not be matched with a Vport on the egress port, the host queue will not be bandwidth controlled and will compete for bandwidth directly based on its own PIR and CIR parameters.

By default, a subscriber host queue with the port-parent option enabled is scheduled within the context of the port's port scheduler policy.

Parameters *description-string* — The destination character string. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

mon-port-sch

Syntax **[no] mon-port-sch**

Context config>port>ethernet
config>port>ethernet>access>egress>vport
config>port>sonet-sdh>path>access>egress>vport

Description This command enables congestion monitoring on an Egress Port Scheduler (EPS) that is applied to a physical port or to a Vport.

Congestion monitoring must be further configured under the port-scheduler CLI hierarchy. Once the congestion monitoring is in effect, the offered rate (incoming traffic) is compared to the configured port-scheduler congestion threshold. The results of these measurements are stored as the number of samples representing the number of times the offered rates exceeded the configured congestion threshold since the last clearing of the stats. Therefore, the results represent the number of times that the port-scheduler that is applied to a port/Vport was congested since the last reset of the stats (via a **clear** command).

The **no** form of the command disables congestion monitoring.

Default no mon-port-sch

port-scheduler-policy

Syntax **port-scheduler-policy** *port-scheduler-policy-name*
no port-scheduler-policy

Context config>port>ethernet>access>egress>vport
config>port>sonet-sdh>path>access>egress>vport

Description This command specifies the destination and organization strings to be used for matching subscriber hosts with this Vport.

The parent Vport of a subscriber host queue, which has the port-parent option enabled, is determined by matching the destination string *dest* string associated with the subscriber and the organization string *org* string associated with the subscriber host with the strings defined under a Vport on the port associated with the subscriber.

If a given subscriber host queue does not have the port-parent option enabled, it will be foster-parented to the Vport used by this subscriber and which is based on matching the *dest* string and *org* string. If the subscriber could not be matched with a Vport on the egress port, the host queue will not be bandwidth controlled and will compete for bandwidth directly based on its own PIR and CIR parameters.

By default, a subscriber host queue with the port-parent option enabled is scheduled within the context of the port's port scheduler policy.

The no form of the command removes the port-scheduler-policy-name from the configuration. The **agg-rate** *rate*, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command.

Parameters *port-scheduler-policy-name* — Specifies an existing port-scheduler-policy configured in the config>qos context. The name can be up to 32 characters long.

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>port>ethernet>access>egress>vport
config>port>sonet-sdh>path>access>egress>vport

Description This command configures a scheduler policy for the egress virtual port.

Parameters *scheduler-policy-name* — Specifies the name of the scheduler policy, up to 32 characters.

autonegotiate

Syntax **autonegotiate** [limited]
no autonegotiate

Context config>port>ethernet

Description This command enables speed and duplex autonegotiation on Fast Ethernet ports and enables far-end fault indicator support on Gb ports.

There are three possible settings for autonegotiation:

- “on” or enabled with full port capabilities advertised

- “off” or disabled where there are no autonegotiation advertisements
- “limited” where a single speed/duplex is advertised.

When autonegotiation is enabled on a port, the link attempts to automatically negotiate the link speed and duplex parameters. If autonegotiation is enabled, the configured duplex and speed parameters are ignored.

When autonegotiation is disabled on a port, the port does not attempt to autonegotiate and will only operate at the **speed** and **duplex** settings configured for the port. Note that disabling autonegotiation on Gb ports is not allowed as the IEEE 802.3 specification for Gb Ethernet requires autonegotiation be enabled for far end fault indication.

If the **autonegotiate limited** keyword option is specified the port will auto-negotiate but will only advertise a specific speed and duplex. The speed and duplex advertised are the **speed** and **duplex** settings configured for the port. One use for limited mode is for multi-speed Gb ports to force Gb operation while keeping autonegotiation enabled for compliance with IEEE 801.3.

Router requires that autonegotiation be disabled or limited for ports in a Link Aggregation Group to guarantee a specific port speed.

The **no** form of this command disables autonegotiation on this port.

Default	autonegotiate
Parameters	limited — The Ethernet interface will automatically negotiate link parameters with the far end, but will only advertise the speed and duplex mode specified by the Ethernet speed and duplex commands.

dampening

Syntax	dampening
Context	config>port>ethernet
Description	This command enters the context to configure exponential port dampening for an Ethernet port.

Exponential Port Dampening (EPD) reduces the number of physical link transitions reported to upper layer protocols, potentially reducing upper layer protocol churn caused by a faulty link. Penalties are added against a port whenever the port's physical link state transitions from a link-up state to a link-down state. When the penalties exceed a configurable threshold, port-up and -down transitions are no longer advertised to upper layers and the port's operational state will remain down until the penalty amount drops below a configurable reuse threshold. Each transition of link-up state to link-down state increments the accumulated penalty value by 1000. The accumulated penalties for a port are reduced at an exponential decay rate according to a configurable half-life parameter.

half-life

Syntax	half-life <i>half-life</i> max-suppress-time <i>max-time</i>
Context	config>port>ethernet>dampening
Description	This command configures the half-life decay time and the maximum period of time for which the port-up state can be suppressed. The <i>half-life</i> and <i>max-time</i> values must be set at the same time and the ratio of <i>max-time/half-life</i> must be less than or equal to 49 and greater than or equal to 1.
Parameters	<i>half-life</i> — Specifies the time, in seconds, that must pass before penalties decay to one-half the initial amount. Values 1 to 2000 Default 5 <i>max-time</i> — Specifies the maximum suppression time, in seconds, which is the time it can take after the physical link comes up before the worst case accumulated penalties have decayed to the reuse threshold. The maximum penalty is derived from the maximum suppression time, half-life, and reuse threshold, using the following equation: $\text{maximum penalty} = (\text{reuse threshold}) \times 2^{\text{expo:}(\text{max-time/half-life})}$ Values 1 to 43200 Default 20

suppress-threshold

Syntax	suppress-threshold <i>suppress-penalties</i> reuse-threshold <i>reuse-penalties</i>
Context	config>port>ethernet>dampening
Description	This command configures the penalties thresholds at which the port state events to the upper layer are to be dampened (suppress threshold) and then permitted (reuse threshold).
Parameters	<i>suppress-penalties</i> — Specifies the threshold at which the port-up state is suppressed until the accumulated penalties drop below the reuse threshold again. Values 1 to 20000 Default 2000 <i>reuse-penalties</i> — Specifies the threshold at which the port-up state is no longer suppressed, after the port has been in a suppressed state and the accumulated penalties decay drops below this threshold. The reuse threshold value must be less than the suppress threshold value. Values 1 to 20000 Default 1000

dot1q-etype

Syntax	dot1q-etype <i>value</i> no dot1q-etype				
Context	config>port>ethernet				
Description	This command specifies the Ethertype expected when the port's encapsulation type is dot1q. Dot1q encapsulation is supported only on Ethernet interfaces. The no form of this command reverts to the default value.				
Parameters	<i>value</i> — Specifies the Ethertype to expect, in either decimal or hex. <table> <tr> <td>Values</td><td>1536 to 65535 (0x0600 to 0xffff)</td></tr> <tr> <td>Default</td><td>If the encap-type is dot1p, then the default is 0x8100. If the encap-type is qinq, then the default is 0x8100.</td></tr> </table>	Values	1536 to 65535 (0x0600 to 0xffff)	Default	If the encap-type is dot1p, then the default is 0x8100. If the encap-type is qinq, then the default is 0x8100.
Values	1536 to 65535 (0x0600 to 0xffff)				
Default	If the encap-type is dot1p, then the default is 0x8100. If the encap-type is qinq, then the default is 0x8100.				

duplex

Syntax	duplex { full half }
Context	config>port>ethernet
Description	This command configures the duplex of a Fast Ethernet port when autonegotiation is disabled. This configuration command allows for the configuration of the duplex mode of a Fast Ethernet port. If the port is configured to autonegotiate this parameter is ignored.
Default	duplex full
Parameters	full — Sets the link to full duplex mode. half — Sets the link to half duplex mode.

efm-oam

Syntax	efm-oam
Context	config>port>ethernet
Description	This command configures EFM-OAM attributes.

accept-remote-loopback

Syntax	[no] accept-remote-loopback

Context	config>port>ethernet>efm-oam
Description	This command enables reactions to loopback control OAM PDUs from peers. The no form of this command disables reactions to loopback control OAM PDUs.
Default	no accept-remote-loopback

dying-gasp-tx-on-reset

Syntax	[no] dying-gasp-tx-on-reset
Context	config>system>ethernet>efm-oam config>port>ethernet>efm-oam
Description	This command enables generation of the Information OAM PDU off-cycle when the soft reset notification is received by the EFM application. The local port state remains under the control of the Soft Reset application and does not change based on this EFM function. If the port is operationally up then the local node will continue to consider the port as available for service data and forwarding. If the upstream node requires notification to route around the local node undergoing the soft reset, notification must be sent to those nodes. This is a disruptive function. This command is disabled by default at the system level and enabled by default at the port level. The combination of the system-level and port-level configuration determines if the dying gasp on soft reset function is active on individual ports. Both the system-level and port-level commands must be enabled in order to support generation of the Information OAM PDU for soft reset. If either is disabled, dying gasp is not active on those ports. This functionality must be enabled prior to the soft reset. When both grace-tx-enable and dying-gasp-tx-on-reset are active on the same port, grace-tx-enable takes precedence when a soft reset is invoked if the Peer Vendor OUI being received is 00:16:4d (ALU) or the configured grace-vendor-oui value. The grace-tx-enable command should not be configured if the Nokia Vendor Specific Grace TLV is not supported on the remote peer, including Nokia 7750 SR equipment prior to release 11.0 R4.
Default	config>system>ethernet>efm-oam>no dying-gasp-tx-on-reset config>port>ethernet>efm-oam>dying-gasp-tx-on-reset

discovery

Syntax	discovery
Context	config>port>ethernet>efm-oam
Description	This is the top level of the hierarchy containing various discovery parameters that allow the operator to control certain aspects of the negotiation process as well as what action to take when there is a mismatch in peer capabilities.

advertise-capabilities

Syntax	advertise-capabilities
Context	config>port>ethernet>efm-oam>discovery
Description	This is the top level of the hierarchy which allows for the overriding of default advertising of capabilities to a remote peer.

link-monitoring

Syntax	[no] link-monitoring
Context	config>port>ethernet>efm-oam>discovery>advertise-capabilities
Description	When the link monitoring function is in a no shutdown state, the Link Monitoring capability (EV) is advertised to the peer through the EFM OAM protocol. This may not be desired if the remote peer does not support the Link Monitoring functionality. The no version of this command suppresses the advertisement of capabilities
Default	link-monitoring

grace-tx-enable

Syntax	[no] grace-tx-enable
Context	config>system>ethernet>efm-oam config>port>ethernet>efm-oam
Description	Enables the sending of grace for all the enabled EFM-OAM sessions on the node. Disabled by default at the system level and enabled by default at the port level. The combination of the system level and port level configuration will determine if the grace function is enabled on the individual ports. Both the system level and the port level must be enabled in order to support grace on a specific port. If either level is disabled, grace is not enabled on those ports. Enabling grace during an active ISSU or soft reset does not invoke the grace function for the active event. When both grace-tx-enable and dying-gasp-tx-on-reset are active on the same port, grace-tx-enable takes precedence when a soft reset is invoked if the Peer Vendor OUI being received is 00:16:4d (ALU) or the configured grace-vendor-oui value. The grace-tx-enable command should not be configured if the Nokia Vendor Specific Grace TLV is not supported on the remote peer, including Nokia 7750 SR equipment prior to release 11.0 R4. The no form of the command disables the sending of the Nokia Vendor Specific Grace TLV.
Default	config>system>ethernet>efm-oam>no grace-tx-enable config>port>ethernet>efm-oam>grace-tx-enable

grace-vendor-oui

Syntax	grace-vendor-oui <i>oui</i> no grace-vendor-oui
Context	config>port>ethernet>efm-oam
Description	This optional command configures an additional peer vendor OUI which indicates support for the Vendor Specific EFM-OAM Grace functionality, allowing grace to be preferred over dying gasp when both are configured. This is in addition to the Nokia Vendor OUI 00:16:4d. When both grace-tx-enable and dying-gasp-tx-on-reset are active on the same port, grace-tx-enable takes precedence when a soft reset is invoked if the Peer Vendor OUI being received is 00:16:4d (ALU) or the configured grace-vendor-oui value. The grace-tx-enable command should not be configured if the Nokia Vendor Specific Grace TLV is not supported on the remote peer, including Nokia 7750 SR equipment prior to release 11.0 R4. The no form of the command removes the additional Vendor OUI but does not remove the Nokia 00:16:4d value.
Default	no grace-vendor-oui
Parameters	<i>oui</i> — Hex value in the range 00:00:00 to FF:FF:FF.

hold-time

Syntax	hold-time <i>time-value</i> no hold-time
Context	config>port>ethernet>efm-oam
Description	This command configures efm-oam operational transition dampening timers which reduce the number of efm-oam state transitions reported to upper layers.
Default	no hold-time
Parameters	<i>time-value</i> — Indicates the number of seconds that the efm-oam protocol will wait before going back to the operational state after leaving the operational state. Note that the hold-time does not apply if efm-oam moved from operational to link-fault. A hold-time value of zero indicates that there should be no delay in transitioning to the operational state. A non-zero value will cause the efm-oam protocol to attempt to negotiate with a peer if possible, but it will remain in the send-local-remote-ok state until the hold time has expired if negotiation is successful. If efm-oam is administratively shutdown while it was in the operational state and then re-enabled when a non-zero hold time is configured, efm-oam will attempt transition to the operational state immediately.

The **no** form of this command reverts the value to the default.

Values 0 to 50

Default 0

ignore-efm-state

Syntax	[no] ignore-efm-state
Context	config>port>ethernet>efm-oam>
Description	When the ignore-efm-state command is configured, any failure in the protocol state machine (discovery, configuration, timeout, loops, and so on) does not impact the state of the port. There is only be a protocol warning message on the port. If this optional command is not configured, the port state is affected by any existing EFM-OAM protocol fault condition.
Default	no ignore-efm-state

link-monitoring

Syntax	link-monitoring
Context	config>port>ethernet>efm-oam
Description	This context contains link monitoring specific options defining the various local thresholds, port interaction and peer notification methods. In order to activate Link monitoring function, this context must be configured with the no shutdown option. Shutting down link monitoring will clear all historical link monitoring counters. If the port was removed from service and placed in a non-operational down state and a port state of link up because a signal failure threshold was crossed and link monitoring is shutdown, the port will be returned to service assuming no underlying conditions prevent this return to service. When the link monitoring function is in a no shutdown state, the Link Monitoring capability (EV) is advertised to the peer through the EFM OAM protocol. This may not be desired if the remote peer does not support the Link Monitoring functionality.

errored-frame

Syntax	errored-frame
Context	config>port>ethernet>efm-oam>link-monitoring
Description	The context used to define errored frame parameters including thresholds, and windows of time to which the error count will be compared. An errored frame is counted when there is any frame error detected by the Ethernet physical layer. This excludes jumbo frames above 9192 bytes which are dropped prior to this function.

event-notification

Syntax	[no] event-notification
Context	config>port>ethernet>efm-oam>link-mon>errored-frame config>port>ethernet>efm-oam>link-mon>errored-frame-period config>port>ethernet>efm-oam>link-mon>errored-frame-seconds
Description	Allows the frame error sf-threshold crossing events to transmit the Event Notification OAMPDU with the specific Link Event TLV information. The Event Notification OAM PDU will only be generated when the initial sf-threshold is reached. No subsequent notification will be sent until the event that triggered until the event is manually cleared. The burst parameter under the local-sf-action will determine the number of Event Notification OAMPDUs to generate when the event occurs. The reception of the event notification will be processed regardless of this parameter. The no version of this command will disable the transmission of the Event Notification OAMPDU for this event type.
Default	event-notification

sd-threshold

Syntax	sd-threshold <i>errored-frames</i> no sd-threshold
Context	config>port>ethernet>efm-oam>link-mon>errored-frame
Description	The option defines the number of errored frames within the configured window which indicates the port has gone beyond an acceptable error rate and should be considered degraded. This is a first level warning that a port may be suspect. This generates an information log event message only and will be recorded in the Port event index but has no port level actions when the error count is equal to or greater than the threshold. This value must be lower than or equal to the sf-threshold value. The no value of this option disables the sd-threshold.
Default	no sd-threshold
Parameters	<i>errored-frames</i> — Specifies the number of errored frames within the configured window which indicates the port has become degraded.
Values	1 to 1000000

sf-threshold

Syntax	sf-threshold <i>errored-frames</i>

Context	config>port>ethernet>efm-oam>link-mon>errored-frame
Description	The option defines the number of frame errors within the configured window which indicates the port has exceeded an acceptable error rate. A log event will be raised, and the port will be taken out of service by default. Configuration options exist to take additional actions when the error rate exceeds the threshold. These actions are defined using the local-sf-action configuration. This event can only be cleared through manual intervention that affects the state of the port.
Parameters	<i>errored-frames</i> — The number of errored frames within the configured window which indicates the port has become unusable.
Values	1 to 1000000
Default	1

window

Syntax	window <i>deciseconds</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-frame
Description	This command defines the size of the window using a 100ms base <i>deciseconds</i> . Errors are accumulated until the end of the window. At the end of the window the actual errors are compared to the thresholds to determine if a threshold has been crossed. There is no mid-window threshold checking. The window represents a unique non-overlapping period of time.
Default	window 10
Parameters	<i>deciseconds</i> — The number of 100ms increments. Must be specified in increments of 10 (full seconds).
Values	10 to 600

errored-frame-period

Syntax	errored-frame-period
Context	config>port>ethernet>efm-oam>link-monitoring
Description	The context used to define errored frame parameters including thresholds, and windows of received packets to which the error count will be compared. An errored frame is counted when there is any frame error detected by the Ethernet physical layer. This excludes jumbo frames above 9192 bytes which are dropped prior to this function. The received packet count will be check every one second to see if the window has been reached.

sd-threshold

Syntax	sd-threshold <i>errored-frames</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-frame-period
Description	The option defines the number of errored frames within the configured window which indicates the port has gone beyond an acceptable error rate and should be considered degraded. This is a first level warning that a port may be suspect. This generates an information log event message only and will be recorded in the Port event index but has no port level actions when the error count is equal to or greater than the threshold. This value must be lower than or equal to the sf-threshold value. The no value of this option disables the sd-threshold
Default	no sd-threshold
Parameters	<i>errored-frames</i> — Specifies the number of errored frames within the configured window which indicates the port has become degraded.
Values	1 to 1000000

sf-threshold

Syntax	sf-threshold <i>errored-frames</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-frame-period
Description	This command defines the number of frame errors within the configured window which indicates the port has exceeded an acceptable error rate. A log event will be raised, and the port will be taken out of service by default. Configuration options exist to take additional actions when the error rate exceeds the threshold. These actions are defined using the local-sf-action configuration. This event can only be cleared through manual intervention that affects the state of the port.
Default	sf-threshold 1
Parameters	<i>errored-frames</i> — Specifies the number of errored frames within the configured window which indicates the port has become unusable.
Values	1 to 1000000

window

Syntax	window <i>packets</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-frame-period

Description	This command defines the size of the window based on a packet receive rate. The minimum serviceable rate is the number of minimum size packets that can be received in one second. The window receive count value will be polled at a minimum one second intervals to see if the window size has been reached. Errors are accumulated until the end of the window. At the end of the window the actual errors are compared to the thresholds to determine if a threshold has been crossed. There is no mid-window threshold checking. The window represents a unique non-overlapping period of time.
Parameters	<i>packets</i> — Specifies the number of received packets in one second.
Values	1 to 4294967295
Default	1488095

errored-frame-seconds

Syntax	errored-frame-seconds
Context	config>port>ethernet>efm-oam>link-monitoring
Description	This command defines the errored frame seconds parameters including thresholds, and windows of time to which the error count will be compared. An errored second is any second in which a single frame error occurred. An errored frame is counted when there is any frame error detected by the Ethernet physical layer. This excludes jumbo frames above 9192 bytes that are dropped prior to this function.

sd-threshold

Syntax	sd-threshold <i>errored-frames</i> no sd-threshold
Context	config>port>ethernet>efm-oam>link-mon>errored-frame-seconds
Description	This command defines the number of errored frame seconds within the configured window which indicates the port has gone beyond an acceptable error rate and should be considered degraded. This is a first level warning that a port may be suspect. This event is raised when the error count is equal to or greater than the configured threshold. This is an information log event message only and will be recorded in the Port event index but has no port level actions. This value must be lower than or equal to the sf-threshold value. The no version of this command disables the sd-threshold.
Parameters	<i>errored-frames</i> — Specifies the number of errored seconds within the configured window which indicates the port has become degraded.
Values	1 to 900

sf-threshold

Syntax	sf-threshold <i>errored-seconds</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-frame-seconds
Description	This command defines the number of errors seconds within the configured window which indicates the port has exceeded an acceptable error rate. A log event will be raised, and the port will be taken out of service by default. Configuration options exist to take additional actions when the error rate exceeds the threshold. These actions are defined using the local-sf-action configuration. This event can only be cleared through manual intervention that affects the state of the port.
Parameters	<i>errored-seconds</i> — Specifies the number of errored seconds within the configured window which indicates the port has become unusable. Values 1 to 900

window

Syntax	window <i>deciseconds</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-frame-seconds
Description	This command defines the size of the window using a 100ms base <i>deciseconds</i> . Errored seconds are accumulated until the end of the window. At the end of the window, the actual errors are compared to the thresholds to determine if a threshold has been crossed. There is no mid-window threshold checking. The window represents a unique non-overlapping period of time.
Parameters	<i>deciseconds</i> — Specifies the number of 100 ms increments. Must be specified in increments of 10 (full seconds). Values 100 to 9000 Default 600

errored-symbols

Syntax	errored-symbols
Context	config>port>ethernet>efm-oam>link-monitoring
Description	This command defines the symbol error parameters including thresholds, and windows of time (converted to symbols in that time) to which the error count will be compared. A symbol error occurs when any encoded symbol is in error and independent of frame counters.

event-notification

Syntax	[no] event-notification
Context	config>port>ethernet>efm-oam>link-mon>errored-symbols
Description	This command allows the symbol error event threshold crossing actions to transmit the Event Notification OAM PDU with the specific Link Event TLV information. The Event Notification OAM PDU will only be generated on the initial sf-threshold is reached. No subsequent notification will be sent until the event that triggered the notification clears, through manual intervention or a window where the configured sd-threshold is not reached. The burst parameter under the local-sf-action will determine the number of Event Notification OAM PDUs to generate when the event occurs. The reception of the event notification will be processed regardless of this parameter. The no version of this command will disable the transmission of the Event Notification OAM PDU for this event type.
Default	event-notification

sd-threshold

Syntax	sd-threshold <i>errored-symbols</i> no sd-threshold
Context	config>port>ethernet>efm-oam>link-mon>errored-symbols
Description	This command defines the number of errored frames within the configured window which indicates the port has gone beyond an acceptable error rate and should be considered degraded. This is a first level warning that a port may be suspect. An event is raised when the error count is equal to or greater than this value. This is an information log event message only and will be recorded in the Port event index but has no port level actions. This value must be lower than or equal to the sf-threshold value. Specific to symbol errors, this value must be configured with the value that indicates anything less is acceptable and the port can be returned to service. If this value is not configured then manual operation is required to return the port to service. The no value of this option means there is there is no automatic return to service.
Default	no sd-threshold
Parameters	<i>errored-symbols</i> — Specifies the number of errored symbols which indicates the port has become degraded. Values 1 to 1000000

sf-threshold

Syntax	sf-threshold <i>errored-symbols</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-symbols
Description	This command defines the number of symbol errors within the configured window which indicates the port has exceeded an acceptable error rate. A log event will be raised, and the port will be taken out of service by default. Configuration options exist to take additional actions when the error rate exceeds the threshold. These actions are defined using the local-sf-action configuration.
Parameters	<i>errored-symbols</i> — Specifies the number of errored-symbols which indicates the port has become unusable. Values 1 to 1000000

window

Syntax	window <i>deciseconds</i>
Context	config>port>ethernet>efm-oam>link-mon>errored-symbols
Description	This command defines the size of the window using a 100ms base <i>deciseconds</i> . The time value is converted to a number of symbols for the underlying medium. Errors are accumulated until the end of the window. At the end of the window, the actual errors are compared to the thresholds to determine if a threshold has been crossed. There is no mid-window threshold checking. The window represents a unique non-overlapping period of time.
Parameters	<i>deciseconds</i> — Specifies the number of 100ms increments in increments of 10 (full seconds). Values 10 to 600 Default 10

shutdown

Syntax	[no] shutdown
Context	config>port>ethernet>efm-oam>link-monitoring
Description	This command enables the link monitoring function. Issuing a no shutdown will start the process. Issuing a shutdown will clear any previously established negative conditions that were a result of the link monitoring process on this port and all collected data. This also controls the advertising capabilities. The no form of the command activates the link monitoring function.
Default	shutdown

shutdown

Syntax	[no] shutdown
Context	config>port>ethernet>efm-oam>link-mon>errored-frame config>port>ethernet>efm-oam>link-mon>errored-frame-period config>port>ethernet>efm-oam>link-mon>errored-frame-seconds config>port>ethernet>efm-oam>link-mon>errored-symbols
Description	This command enables or disables the local counting, thresholding and actions associated with this type of local monitor. Peer received errors are not controlled by this command. Reaction to peer messaging is defined in the peer-rdi-rx hierarchy. The no form of the command activates the local monitoring function and actions for the event.
Default	shutdown

local-sf-action

Syntax	local-sf-action
Context	config>port>ethernet>efm-oam>link-monitoring
Description	This command defines how crossing the local signal failure threshold (sf-threshold) will be handled. This includes local actions and if and how to notify the peer that the threshold has been crossed.

event-notification-burst

Syntax	event-notification-burst <i>packets</i>
Context	config>port>ethernet>efm-oam>link-mon>local-sf-action
Description	This command defines the number of the Event Notification OAM PDU to be send to the peer if the local signal failure threshold (sf-threshold) has been reached. The sending of the Event Notification OAMPDU is configured under the individual monitors. Interactions: The sf-thresh threshold will trigger these actions.
Parameters	<i>packets</i> — Specifies the number of Event Notification OAM PDUs to send to a peer when the signal failure threshold has been reached.
Values	1 to 5

info-notification

Syntax	info-notification

Context config>port>ethernet>efm-oam>link-mon>local-sf-action

Description The context allows the operator to set different flags in the Information OAM PDU. The flags can be used to notify the peer that a local signal failure threshold has been exceeded within the configured window. This is useful when the local node supports the link monitoring function, but the remote peer does not support this capability. Information OAM PDUs are sent on the interval where the Event Notification OAM PDU is typically only sent on the initial sf-threshold crossing event. It is strongly suggested one of the Information OAM PDU Flag fields used to continually communicate current monitor state to the peer.

Interactions: The signal failure threshold will trigger these actions.

critical-event

Syntax [no] critical-event

Context config>port>ethernet>efm-oam>link-mon>local-sf-action>info-notification

Description This command sets the critical event Flag field in the Information OAMPDU when the local signal failure (sf-threshold) threshold is reached. This will be maintained in all subsequent Information OAM PDUs until the situation is cleared.

Interactions: The signal failure threshold will trigger these actions.

Default no critical-event

dying-gasp

Syntax [no] dying-gasp

Context config>port>ethernet>efm-oam>link-mon>local-sf-action>info-notification

Description This command sets the dying gasp Flag field in the Information OAM PDU when the local signal failure (sf-threshold) threshold is reached. This will be maintained in all subsequent Information OAM PDUs until the situation is cleared.

Interactions: The signal failure threshold will trigger these actions.

Default no dying-gasp

local-port-action

Syntax local-port-action {log-only | out-of-service}

Context config>port>ethernet>efm-oam>link-mon>local-sf-action

Description	This command configures the parameters that define if and how the local port will be affected when the local signal failure threshold (sf-threshold) has been reached within the configured window. Interactions: The signal failure threshold will trigger these actions.
Default	local-port-action out-of-service
Parameters	log-only — Keyword that prevents the port from being affected when the configured signal failure threshold is reach within the window. The event will be logged but the port will remain operational. out-of-service — Keyword that causes the port to enter a non-operation down state with a port state of link up. The error will be logged when the configured signal failure threshold (sf-threshold) is reached within the window. The port will not be available to service data but will continue to carry Link OAM traffic to ensure the link is monitored.

mode

Syntax	mode {active passive}
Context	config>port>ethernet>efm-oam
Description	This command configures the mode of OAM operation for this Ethernet port. These two modes differ in that active mode causes the port to continually send out efm-oam info PDUs while passive mode waits for the peer to initiate the negotiation process. A passive mode port cannot initiate monitoring activities (such as loopback) with the peer.
Default	mode active
Parameters	active — Provides capability to initiate negotiation and monitoring activities. passive — Relies on peer to initiate negotiation and monitoring activities.

peer-rdi-rx

Syntax	peer-rdi-rx
Context	config>port>ethernet>efm-oam
Description	This container allows an action to be configured for the various event conditions that can be received from a peer under the context of the EFM OAM protocol.

critical-event

Syntax	critical-event local-port-action {log-only out-of-service}

Context	config>port>ethernet>efm-oam>peer-rdi-rx
Description	This command defines how to react to the reception of a critical event Flag field set in the informational OAMPDU.
Default	critical-event local-port-action out-of-service
Parameters	local-port-action — Defines whether or not the local port will be affected when a critical event is received from a peer. log-only — Keyword that prevents the port from being affected when the local peer receives a critical event. The critical event will be logged but the port will remain operational. out-of-service — Keyword that causes the port to enter a non-operation down state with a port state of link up. The error will be logged upon reception of critical event. The port will not be available to service data but will continue to carry Link OAM traffic to ensure the link is monitored.

dying-gasp

Syntax	dying-gasp local-port-action {log-only out-of-service}
Context	config>port>ethernet>efm-oam>peer-rdi-rx
Description	This command defines how to react to the reception of a dying gasp Flag field set in the informational OAMPDU.
Default	dying-gasp local-port-action out-of-service
Parameters	local-port-action — Defines whether or not the local port will be affected when a dying gasp event is received from a peer. log-only — Keyword that prevents the port from being affected when the local peer receives a dying gasp. The dying gasp will be logged but the port will remain operational. out-of-service — Keyword that causes the port to enter a non-operation down state with a port state of link up. The error will be logged upon reception of dying gasp. The port will not be available to service data but will continue to carry Link OAM traffic to ensure the link is monitored.

event-notification

Syntax	event-notification local-port-action {log-only out-of-service}
Context	config>port>ethernet>efm-oam>peer-rdi-rx

Description	This command defines how to react to the reception of event TLVs contained in the Event Notification OAMPDU. The event TLVs contained in the event notification OAMPDU will be analyzed to determine if the peer has crossed the error threshold for the window. The analysis does not consider any local signal degrades or signal failure threshold. The analysis is based solely on the information receive form the peer. The analysis is performed on all event TLVs contained in the Event Notification OAMPDU without regard for support of a specific error counters or local configuration of any thresholds. In the case of symbol errors only, a threshold below the error rate can be used to return the port to service.
Default	event-notification local-port-action log-only
Parameters	local-port-action — Defines whether or not the local port will be affected when the Event Notification OAM PDU is received from a peer based on the threshold computation for the included TLVs. log-only — Keyword that prevents the port from being affected when the local peer receives an Event Notification OAM PDU. The event will be logged but the port will remain operational. out-of-service — Keyword that causes the port to enter a non-operation down state with a port state of link up. The error will be logged upon reception of Event Notification. The port will not be available to service data but will continue to carry Link OAM traffic to ensure the link is monitored. All this assumes the error threshold exceeds the error rate in the TLV.

link-fault

Syntax	link-fault local-port-action {log-only out-of-service}
Context	config>port>ethernet>efm-oam>peer-rdi-rx
Description	This command defines how to react to the reception of a link fault flag set in the informational PDU from a peer.
Default	link-fault local-port-action out-of-service
Parameters	local-port-action — Defines whether or not the local port will be affected when a link fault is received from a peer. log-only — Keyword that prevents the port from being affected when the local peer receives a link fault. The dying gasp will be logged but the port will remain operational. out-of-service — Keyword that causes the port to enter a non-operation down state with a port state of link up. The error will be logged upon reception of link fault event. The port will not be available to service data but will continue to carry Link OAM traffic to ensure the link is monitored.

transmit-interval

Syntax	[no] transmit-interval <i>interval</i> [multiplier <i>multiplier</i>]
Context	config>port>ethernet>efm-oam
Description	This command configures the transmit interval of OAM PDUs.
Default	transmit-interval 10 multiplier 5
Parameters	<i>interval</i> — Specifies the transmit interval, in 100 milliseconds. Values 1 to 600 <i>multiplier</i> — Specifies the multiplier for transmit-interval to set local link down timer. Values 2 to 5

trigger-fault

Syntax	trigger-fault {dying-gasp critical-event} no trigger-fault
Context	config>port>ethernet>efm-oam
Description	This command configures the appropriate flag field in the Information OAM PDU, bursting three consecutive packets during the off cycle. If the local port state is operational, this command changes the local port state to “Link Up”. If the local port state is not operational, this configuration is installed as an EFM reason to prevent the port from returning to an Up operational state. This command can be used as a precursor to a port shutdown. This terminates the peering relationship without having to wait for protocol timeouts, assuming the peer supports the necessary action when receiving the dying gasp or critical event flag setting. The no form of this command disables this functionality.
Default	no trigger-fault
Parameters	dying-gasp — Keyword to set the dying gasp flag. critical-event — Keyword to set the critical event flag.

tunneling

Syntax	[no] tunneling
Context	config>port>ethernet>efm-oam
Description	This command enables EFM OAM PDU tunneling. Enabling tunneling will allow a port mode Epipe SAP to pass OAM frames through the pipe to the far end.

The **no** form of the command disables tunneling.

Default no tunneling

egress-rate

Syntax **egress-rate** *sub-rate*
no egress-rate

Context config>port>ethernet

Description This command configures the rate of traffic leaving the network. The configured *sub-rate* uses packet-based accounting. An event log is generated each time the egress rate is modified unless the port is part of a LAG.

This command is not supported on ports of the following MDA types:

- m60-10/100eth-tx
- c8-10/100eth-tx
- m10-1gb-hs-sfp-b
- m1-10gb-hs-xfp-b

The **no** form of this command returns the value to the default.

Default no egress-rate

Parameters *sub-rate* — Specifies the egress rate in kb/s.

Values 1 to 100000000

encap-type

Syntax **encap-type** {**dot1q** | **null** | **qinq**}
no encap-type

Context config>port>ethernet

Description This command configures the encapsulation method used to distinguish customer traffic on an Ethernet access port, or different VLANs on a network port.

The **no** form of this command restores the default.

Default encap-type null

Parameters **dot1q** — Ingress frames carry 802.1Q tags where each tag signifies a different service.
null — Ingress frames will not use any tags to delineate a service. As a result, only one service can be configured on a port with a null encapsulation type.

qinq — Specifies QinQ encapsulation.

eth-bn-egress-rate-changes

Syntax	eth-bn-egress-rate-changes no eth-bn-egress-rate-changes
Context	config>port>ethernet
Description	This command allows rate changes received in ETH-BN messages on a port-based MEP to update the egress rate used on the port. The egress rate is capped by the minimum of the configured egress-rate and the maximum port rate, and the minimum egress rate is 1 kb/s. The no form of this command returns the value to the default. This command is not supported on the following MDA types: • m60-10/100eth-tx • c8-10/100eth-tx • m10-1gb-hs-sfp-b • m1-10gb-hs-xfp-b
Default	no eth-bn-egress-rate-changes

hold-time

Syntax	hold-time {[up <i>hold-time-up</i>] [down <i>hold-time-down</i>] [seconds centiseconds]}
	no hold-time
Context	config>port>ethernet
Description	This command configures port link dampening timers which reduce the number of link transitions reported to upper layer protocols. The hold-time value dampens interface transitions. When an interface transitions from an up state to a down state, it is immediately advertised to the rest of the system if the hold-time down interval is zero, but if the hold-time down interval is greater than zero, interface down transitions are not advertised to upper layers until the hold-time down interval has expired. Likewise, an interface is immediately advertised as up to the rest of the system if the hold-time up interval is zero, but if the hold-time up interval is greater than zero, up transitions are not advertised until the hold-time up interval has expired.

For ESM SRRP setup, MCS synchronizes subscriber information between the two chassis. After a chassis recovers from a power reset/down, MCS immediately synchronizes all subscriber information at once. The longer the host list, the longer it will take to synchronize the chassis. In a fully populated chassis, it is recommended to allow at least 45 minutes for MCS synchronization. It is also recommended to hold the port down, facing the subscriber, on the recovering chassis for 45 minutes before it is allowed to forward traffic again.

The **no** form of this command reverts to the default values.

Default down 0 seconds — No port link down dampening is enabled; link down transitions are immediately reported to upper layer protocols.

up 0 seconds — No port link up dampening is enabled; link up transitions are immediately reported to upper layer protocols.

Parameters *hold-time-up* — The delay, in seconds or centiseconds, to notify the upper layers after an interface transitions from a down state to an up state.

Values 0 to 36000 seconds
0, 10 to 3600000 centiseconds in 5 centisecond increments

hold-time-down — The delay, in seconds or centiseconds, to notify the upper layers after an interface transitions from an up state to a down state.

Values 0 to 36000 seconds
0, 10 to 3600000 centiseconds in 5 centisecond increments

seconds | centiseconds — Specifies the units of your hold time in **seconds** or **centiseconds**.

hsmdda-scheduler-overrides

Syntax [no] hsmdda-scheduler-overrides

Context config>port>ethernet

Description This command enables the context to configure ingress and egress HSMDDA scheduler override parameters. Executing hsmdda-scheduler-override places the current CLI context into the egress scheduler override node either at the ingress MDA or egress port level.

Default values are listed in [Table 48](#).

Table 48 Default Values

Command	Configuration
description	no description
max-rate	no max-rate

Table 48 **Default Values (Continued)**

Command	Configuration
group	group 1 rate max group 2 rate max
scheduling-class	scheduling-class 1 rate max scheduling-class 2 rate max scheduling-class 3 rate max scheduling-class 4 rate max scheduling-class 5 rate max scheduling-class 6 rate max scheduling-class 7 rate max scheduling-class 8 rate max

The **no** form of the command removes the overridden parameters from the HSM DA egress port or ingress MDA scheduler. Once existing overrides are removed, the scheduler reverts all scheduling parameters back to the parameters defined on the hsm da-scheduler-policy associated with the egress port or ingress MDA.

Parameters **create** — Mandatory for creating an entry.

group

Syntax **group** *group-id* **rate** *rate*
no group *group-id*

Context config>port>ethernet>hsm da-scheduler-overrides

Description This command changes the maximum rate allowed for a weighted scheduling group on the local HSM DA scheduler. Scheduling classes within the group are managed with an aggregate rate limit when either an explicit group rate is defined on the HSM DA scheduling policy or a local override is defined based on the group override command.

The **no** form of the command removes the local overrides for the weighted scheduling group. Once removed, the defined behavior within the HSM DA scheduling policy for the weighted scheduling group is used.

Parameters *group-id* — Identifies the two weighted scheduling groups to be overridden.

Values 1, 2

rate — The *megabits-per-second* parameter specifies a local limit on the total bandwidth for the weighted scheduling group and overrides any rate defined in the HSMDB scheduler policy for the weighted scheduling group. The parameter is specified in Megabits per second in a base 10 context. A value of 1 equals a rate of 1000000 bits per second.

The **max** keyword removes any existing rate limit imposed by the HSMDB scheduler policy for the weighted scheduling group allowing it to use as much total bandwidth as possible.

Values 1 to 10000, max (Mb/s)

max-rate

Syntax	max-rate <i>rate</i> no max-rate
Context	config>port>ethernet>hsmdb-scheduler-overrides
Description	This command overrides the max-rate parameters configured in the hsmdb-scheduler-policy associated with the egress port or ingress MDA. When a max-rate is defined at the override level, the HSMDB scheduler policy's max-rate parameter is ignored. The hsmdb-scheduler-override max-rate command supports a max parameter that allows the override command to restore the default of not having a rate limit on the port scheduler. This is helpful when the HSMDB scheduler policy has an explicit maximum rate defined and it is desirable to remove this limit at the port instance. The no form of the command removes the maximum rate override from the egress port or the ingress MDA scheduler context. Once removed, the max-rate parameter from the HSMDB scheduler policy associated with the port or MDA will be used by the local scheduler context.
Parameters	<i>rate</i> — The rate parameter is mutually exclusive to specifying the max keyword. When executing the max-rate override command either the keyword max or a rate in megabits-per-second must be specified. Values 1 to 10000, max (Mb/s) max — The max keyword is mutually exclusive to specifying a rate in megabits-per-second. When executing the max-rate override command either the keyword max or a rate in megabits-per-second must be specified. The max keyword removes an existing rate limit from the HSMDB scheduler context.

scheduling-class

Syntax	scheduling-class <i>class rate rate</i> scheduling-class <i>class weight weight-in-group</i> no scheduling-class <i>class</i>

Context	config>port>ethernet>hsmda-scheduler-overrides
Description	This command overrides the maximum rate allowed for a scheduling class or the weight of the class within a weighted scheduling group. The scheduling-class override cannot be used to change scheduling class weighted group membership; weighted group membership may only be defined within the HSMDA scheduling policy. Scheduling classes correspond directly to the queue-IDs used by every queue on an HSMDA. All queues with an ID of 1 associated with the scheduler are members of scheduling class 1 on the scheduler. Queues with an ID of 2 are members of scheduling class 2. This is true through scheduling class 8. When the scheduling class is not a member of a weighted group, the scheduling-class command may be used to modify the maximum rate allowed for the scheduling class. This is done using the rate parameter followed by either the max keyword or an actual rate defined as megabits-per-second. Use the rate max combination to locally remove a rate limit defined for the class on the scheduling policy. When the rate megabits-per-second combination is used, the scheduling class defined as class-id is rate limited to the specified rate. Either the max keyword or a value for megabits-per-second must follow the rate keyword. The rate keyword is mutually exclusive with the weight keyword. The weight keyword may only be specified when class-id is a member of a weighted scheduling group. When the weight keyword is specified, a weight value specified as weight must follow. The new weight locally overrides the weight defined for the scheduling class in the HSMDA scheduling policy. When the scheduling-class command is executed, either the rate or weight keyword must follow. When a scheduling class has a local rate override, the HSMDA policy associated with the override cannot move the scheduling class into a weighted scheduling group. Similarly, when a scheduling class has a local weight override, the HSMDA policy associated with the override cannot define a rate (neither max nor a megabit-per-second value) for the scheduling class. The local overrides of the scheduling class must be removed before these changes may be made. The no form of the command removes the local overrides for the scheduling class. Once removed, the defined behavior for the scheduling class within the HSMDA scheduling policy will used.
Parameters	<i>class</i> — Identifies the scheduling class to be being overridden. Values 1 to 8 <i>rate</i> — Overrides the HSMDA scheduler policies maximum rate for the scheduling class and requires either the max keyword or a rate defined in megabits-per-second. In order for the rate keyword to be specified, the scheduling class cannot be a member of a weighted scheduling group as defined on the HSMDA scheduling policy. The rate keyword is mutually exclusive with the weight keyword. Also, either the rate or weight keyword must be specified.

The **max** keyword removes any existing rate limit imposed by the HSMDA scheduler policy for the scheduling class allowing it to use as much total bandwidth as possible.

Values 1 to 100000, max (Mb/s)

weight-in-group — Overrides the weighted scheduler group weight for the scheduling class as defined in the HSMDA scheduler policy. In order for the weight keyword to be specified, the scheduling class must be a member of a weighted scheduling group as defined on the HSMDA scheduling policy. A value represented by group-weight must follow the **weight** keyword. The new weight will be used to determine the bandwidth distribution for member scheduling classes within the group of which the scheduling class is a member.

Values 1 to 100

class

Syntax	[no] class <i>class-number</i>
Context	config>port>ethernet>egress>hs-sec-shaper
Description	This command specifies the HS secondary shaper class. The no form of the command reverts the rate for this class to the default value.
Parameters	<i>class-number</i> — Specifies the HS secondary shaper class identifier.
	Values 1 to 6

ingress-rate

Syntax	ingress-rate <i>sub-rate</i> no ingress-rate
Context	config>port>ethernet
Description	This command configures the maximum amount of ingress bandwidth that this port can receive with the configured sub-rate using packet-based accounting. The no form of this command returns the value to the default.
Default	no ingress-rate
Parameters	<i>sub-rate</i> — Specifies the ingress rate, in Mb/s.
	Values 1 to 100000

lacp-tunnel

Syntax	[no] lacp-tunnel
Context	config>port>ethernet
Description	This command enables LACP packet tunneling for the Ethernet port. When tunneling is enabled, the port will not process any LACP packets but will tunnel them instead. The port cannot be added as a member to a LAG group. In this context, the lacp-tunnel command is supported for Epipe and VPLS services only. The no form of the command disables LACP packet tunneling for the Ethernet port.
Default	no lacp-tunnel

load-balancing-algorithm

Syntax	load-balancing-algorithm <i>option</i> no load-balancing-algorithm
Context	config>port>ethernet config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command specifies the load balancing algorithm to be used on this port. In the default mode, no load-balancing-algorithm, the port inherits the global settings. The value is not applicable for ports that do not pass any traffic. The configuration of load-balancing-algorithm at logical port level has three possible values: • include-l4 — Enables inherits system-wide settings including Layer 4 source and destination port value in hashing algorithm. • exclude-l4 — Layer 4 source and destination port value will not be included in hashing. • no load-balancing-algorithm — Inherits system-wide settings. The hashing algorithm addresses finer spraying granularity where many hosts are connected to the network. To address more efficient traffic distribution between network links (forming a LAG group), a hashing algorithm extension takes into account Layer 4 information (src/dst L4-protocol port). The hashing index can be calculated according to the following algorithm: If [(TCP or UDP traffic) & enabled] hash (<TCP/UDP ports>, <IP addresses>)

```

else if (IP traffic)
    hash (<IP addresses>)
else
    hash (<MAC addresses>)
endif

```

This algorithm will be used in all cases where IP information in per-packet hashing is included (see LAG and ECMP Hashing). However the Layer 4 information (TCP/UDP ports) will not be used in the following cases:

- Fragmented packets

Default no load-balancing-algorithm

Parameters *option* — Specifies the load balancing algorithm to be used on this port.

Values **include-l4** — Specifies that the source and destination ports are used in the hashing algorithm.
exclude-l4 — Specifies that the source and destination ports are not used in the hashing algorithm.

min-frame-length

Syntax **min-frame-length** *byte-length*

Context config>port>ethernet

Description This command configures the minimum transmitted frame length.

Parameters *byte-length* — Specifies the number of bytes for the minimum frame length.

Values 64, 68

Default 64

pbb-etype

Syntax **pbb-etype** [*ethertype-value*]
no pbb-etype

Context config>port>ethernet

Description This command configures the Ethertype used for PBB encapsulation.

Default no pbb-etype

Parameters *ethertype-value* — Specifies the Ethertype value in the form of 0x600 to 0xffff.

Values 1536 to 65535 (accepted in decimal or hex)

ptp-asymmetry

Syntax **ptp-asymmetry** *nanoseconds*
no ptp-asymmetry

Context config>port>ethernet

Description This command configures the PTP asymmetry delay delta on an Ethernet port. The command corrects for known asymmetry for time of day or phase recovery of PTP packets on both local and downstream PTP clocks.

Default no ptp-asymmetry

Parameters *nanoseconds* — Specifies the value, in nanoseconds, that the forward path delay varies from the mean path delay; the value can be a negative number.

Values -2147483648 to 2147483647

qinq-etype

Syntax **qinq-etype** *qinq-etype-value*
no qinq-etype

Context config>port>ethernet

Description This command configures the Ethertype used for Q-in-Q encapsulation.
The **no** form of this command reverts the qinq-etype value to the default.

Default no qinq-etype

Parameters *qinq-etype-value* — Specifies the qinq-etype to expect in the form of 0x600 to 0xffff.

Values 1536 to 65535 in decimal or hex formats

report-alarm

Syntax [**no**] **report-alarm** [**signal-fail**] [**remote**] [**local**] [**no-frame-lock**] [**high-ber**] [**no-block-lock**] [**no-am-lock**] [**duplicate-lane**]

Context config>port>ethernet

Description This command specifies when and if to generate alarms and alarm clear notifications for this port.

Parameters	signal-fail — Reports an Ethernet signal lost alarm. remote — Reports remote faults. local — Reports local faults. no-frame-lock — Reports a 'not locked on the Ethernet framing sequence' alarm. high-ber — Reports High Bit Error Rate. no-block-lock — Reports 40G/100G PCS Lanes Not Block Locked. no-am-lock — Reports 40G/100G PCS Alignment Marker Loss of Lock. duplicate-lane — Reports 40G/100G PCS Duplicate Lane Marker.

rs-fec-mode

Syntax	rs-fec-mode <i>rs-fec-mode</i> no rs-fec-mode
Context	config>port>ethernet
Description	This command enables RS-FEC on the Ethernet port. RS-FEC Clause 91 is required for QSFP28 and CFP4 optics for short-reach optics.
Default	no rs-fec-mode
Parameters	<i>rs-fec-mode</i> — Specifies the RS-FEC mode to support. Values cl91-514-528

sflow

Syntax	[no] sflow
Context	config>port>ethernet
Description	This command enables sFlow data collection for a port and its SAPs that support sFlow data collection. The no form of this of this command disables sFlow.
Default	no sflow

single-fiber

Syntax	[no] single-fiber
Context	config>port>ethernet

Description	This command enables packet gathering and redirection of IP packets from a single fiber (RX) port of the Ethernet or SONET/SDH interface and redistributes packets to other interfaces through either static routes or policy-based forwarding. This parameter can be applied in conjunction with the strip-label command. If they are applied together, the port must have the single-fiber option configured before it can be associated with an interface that is configured with the strip-label option. Once a port is configured with single-fiber, traffic will no longer be transmitted out of that port.
Default	no single-fiber

speed

Syntax	speed {10 100 1000 10000 25000 40000 50000 100000}
Context	config>port>ethernet
Description	For ports that support multiple speeds, this command configures the port speed to be used. This applies to the following: • fast Ethernet when autonegotiate is disabled • 10/100/1000 Ethernet when autonegotiate is disabled • 10/1G ports supporting 10G SFP+ or 1G SFP • 40/100G ports supporting QSFP28s on non connector-based MDAs If the port is configured to autonegotiate this parameter is ignored. Speed cannot be configured for ports that are part of a Link Aggregation Group (LAG).
Default	dependent on port type
Parameters	10 — Sets the link to 10 Mb/s speed. 100 — Sets the link to 100 Mb/s speed. 1000 — Sets the link to 1000 Mb/s speed. 10000 — Sets the link to 10000 Mb/s speed. 25000 — Sets the link to 25000 Mb/s speed. 40000 — Sets the link to 40000 Mb/s speed. 50000 — Sets the link to 50000 Mb/s speed. 100000 — Sets the link to 100000 Mb/s speed.

ssm

Syntax	ssm

Context	config>port>ethernet
Description	This command enables the Ethernet Synchronization Messaging Channel (ESMC) for the Ethernet port. ESMC carries the Synchronization Status Message (SSM) code representing the quality level of the source of frequency of the central clock of the node.

code-type

Syntax	code-type [sonet sdh] [no] code-type
Context	config>port>ethernet>ssm
Description	This command configures the encoding of synchronization status messages. For example, whether to use an SDH or SONET set of values. Configuring the network-type is only applicable to SyncE ports. It is not configurable on SONET/SDH ports. For the network-type, sdh refers to ITU-T G.781 Option I, while sonet refers to G.781 Option II (equivalent to Telcordia GR-253-CORE).
Default	code-type sdh
Parameters	sdh — Specifies the values used on a G.781 Option 1 compliant network. sonet — Specifies the values used on a G.781 Option 2 compliant network.

tx-dus

Syntax	[no] tx-dus
Context	config>port>ethernet>ssm config>port>sonet-sdh
Description	This command forces the QL value transmitted from the SSM channel of the SONET/SDH port or the Synchronous Ethernet port to be set to QL-DUS/QL-DNU. This capability is provided to block the use of the interface from the SR/ESS for timing purposes.
Default	no tx-dus

symbol-monitor

Syntax	symbol-monitor
Context	config>port>ethernet

Description This command configures Ethernet Symbol Monitoring parameters. Support for symbol monitoring is hardware dependent. An error message indicating that the port setting cannot be modified will be presented when attempting to enable the feature or configure the individual parameters on unsupported hardware.

sd-threshold

Syntax **sd-threshold** *threshold* [**multiplier** *multiplier*]
no sd-threshold

Context config>port>ethernet>sym-mon

Description This command specifies the error rate at which to declare the Signal Degrade condition on an Ethernet interface. The value represents $M \times 10^{-N}$ a ratio of symbol errors over total symbols received over W seconds of the sliding window. The symbol errors on the interface are sampled once per second. A default of 10 seconds is used when there is no additional window-size configured. The multiplier keyword is optional. If the multiplier keyword is omitted or no sd-threshold is specified the multiplier will return to the default value of 1.

Default no sd-threshold

Parameters *threshold* — Specifies the rate of symbol errors.
Values 1 to 9
multiplier — Specifies the multiplier used to scale the symbol error ratio.
Values 1 to 9

sf-threshold

Syntax **sf-threshold** *threshold* [**multiplier** *multiplier*]
no sf-threshold

Context config>port>ethernet>sym-mon

Description This command specifies the error rate at which to declare the Signal Fail condition on an Ethernet interface. The value represents $M \times 10^{-N}$ symbol errors over total symbols received over W seconds of the sliding window. The symbol errors on the interface are sampled once per second. A default of 10 seconds is used when there is no additional window-size configured. The multiplier keyword is optional. If the multiplier keyword is omitted or no sf-threshold is specified, the multiplier will return to the default value of 1.

Default no sf-threshold

Parameters *threshold* — Specifies the rate of symbol errors.
Values 1 to 9

multiplier — Specifies the multiplier used to scale the symbol error ratio.

Values 1 to 9

window-size

Syntax	window-size <i>seconds</i> no window-size
Context	config>port>ethernet>sym-mon
Description	This command specifies sliding window size over which the symbols are sampled to detect signal failure or signal degraded conditions.
Default	window-size 10
Parameters	<i>seconds</i> — Specifies the size of the sliding window in seconds over which the errors are measured. Values 5 to 60

util-stats-interval

Syntax	util-stats-interval <i>seconds</i>
Context	config>port>ethernet
Description	This command configures the interval used to calculate the utilization statistics. Port utilization statistics are only available for physical Ethernet ports on a host system. These statistics are not available for the following: • Ethernet ports on an Ethernet satellite • Ethernet ports on a VSR • PXC ports • vsm-cca-xp ports
Parameters	<i>seconds</i> — Specifies the size of the interval, in seconds. Values 30 to 600 Default 300

xgig

Syntax	xgig {lan wan}
Context	config>port>ethernet

Description	This command configures a 10 Gb/s interface to be in Local or Wide Area Network (LAN or WAN) mode. When configuring the port to be in WAN mode certain SONET/SDH parameters can be changed to reflect the SONET/SDH requirements for this port. When the port is configured for LAN mode, all SONET/SDH parameters are pre-determined and not configurable.
Default	xgig lan
Parameters	lan — Sets the port to operate in LAN mode. wan — Sets the port to operate in WAN mode.

crc-monitor

Syntax	crc-monitor
Context	config>port>ethernet
Description	This command configures Ethernet CRC Monitoring parameters.

sd-threshold

Syntax	sd-threshold <i>threshold</i> [multiplier <i>multiplier</i>] no sd-threshold
Context	config>port>ethernet>crc-monitor
Description	This command specifies the error rate at which to declare the Signal Degrade condition on an Ethernet interface. The value represents $M \times 10^{-N}$ a ratio of errored frames over total frames received over W seconds of the sliding window. The CRC errors on the interface are sampled once per second. A default of 10 seconds is used when there is no additional window-size configured. The multiplier keyword is optional. If the multiplier keyword is omitted or no sd-threshold is specified the multiplier will return to the default value of 1.
Default	no sd-threshold
Parameters	<i>threshold</i> — Specifies the threshold value. Values 1 to 9 <i>multiplier</i> — Specifies the multiplier value. Values 1 to 9

sf-threshold

Syntax	sf-threshold <i>threshold</i> [multiplier <i>multiplier</i>]

no sf-threshold

Context	config>port>ethernet>crc-monitor
Description	This command specifies the error rate at which to declare the Signal Fail condition on an Ethernet interface. The value represents $M \times 10^E - N$ errored frames over total frames received over W seconds of the sliding window. The CRC errors on the interface are sampled once per second. A default of 10 seconds is used when there is no additional window-size configured. The multiplier keyword is optional. If the multiplier keyword is omitted or no sf-threshold is specified the multiplier will return to the default value of 1.
Default	no sf-threshold
Parameters	<i>threshold</i> — Specifies the threshold value. Values 1 to 9 <i>multiplier</i> — Specifies the multiplier value. Values 1 to 9

window-size

Syntax	window-size <i>seconds</i> no window-size
Context	config>port>ethernet>crc-monitor
Description	This command specifies sliding window size over which the Ethernet frames are sampled to detect signal fail or signal degrade conditions. The command is used jointly with the sf-threshold and the sd-threshold to configure the sliding window size. The no version of this command reverts to the default value of 10 seconds.
Default	no window-size
Parameters	<i>seconds</i> — The size of the sliding window in seconds over which the errors are measured. Values 5 to 60

down-on-internal-error

Syntax	down-on-internal-error [tx-disable] no down-on-internal-error
Context	config>port>ethernet
Description	This command configures the system to bring a port operationally down in the event the system has detected internal MAC transmit errors (Int MAC Tx Errs).

Default	no down-on-internal-error
Parameters	tx-disable — Specifies that the laser should be disabled if an internal MAC transmit error is encountered. When used, this option requires that the operator explicitly cycle the admin state of the port to clear the error and re-enable the laser.

single-fiber

Syntax	[no] single-fiber
Context	config>port>ethernet config>port>sonet-sdh
Description	This command enables packet gathering and redirection of IP packets from a single fiber (RX) port of the Ethernet or SONET/SDH interface and redistributes packets to other interfaces through either static routes or policy-based forwarding. This parameter can be applied in conjunction with the strip-label command. If they are applied together, the port must have the single-fiber option configured before it can be associated with an interface that is configured with the strip-label option. Once a port is configured with single-fiber, traffic will no longer be transmitted out of that port. This command can be used in conjunction with strip-label.
Default	no single-fiber

2.20.2.15 802.1x Port Commands

macsec

Syntax	[no] macsec
Context	config>port>ethernet>dot1x
Description	This command configures macsec under this port.

exclude-mac-policy

Syntax	exclude-mac-policy <i>mac-policy-id</i> no exclude-mac-policy
Context	config>port>ethernet>dot1x>macsec
Description	This command specifies the MAC policy to be excluded from MACsec encryption.

The **no** form of the command removes the policy from the MACsec and allows all destination MAC addresses.

Default	no exclude-mac-policy
Parameters	<i>mac-policy-id</i> — Specifies the MAC policy to exclude from the configuration.
Values	0 to 4294967295

exclude-protocol

Syntax	[no] exclude-protocol { <i>protocol-name</i> }
Context	config>port>ethernet>dot1x>macsec
Description	Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a port. When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link. When this option is enabled on a port where MACsec is configured, packets of the specified protocols will be sent and accepted in clear text.
Default	no exclude-protocol
Parameters	<i>protocol-name</i> — Specifies the protocol name.
Values	cdp, lACP, lldp, eapol-start, efm-oam, eth-cfm, ptp, or ubfd

rx-must-be-encrypted

Syntax	[no] rx-must-be-encrypted
Context	config>port>ethernet>dot1x>macsec
Description	When the rx-must-be-encrypted option is enabled, all traffic that is not MACsec-secured that is received on the port is dropped. When the rx-must-be-encrypted option is disabled, all arriving traffic, whether MACsec secured or not, will be accepted.



Note: This command is only available on the NULL port level and does not have per-VLAN granularity.

The **no** form of this command disables the **rx-must-be encrypted** option.

Default rx-must-be-encrypted

sub-port

Syntax **sub-port** *port-id* [**create**]
no sub-port *port-id*

Context config>port>ethernet>dot1x>macsec

Description This command creates a MACsec instance on a physical port, targeting the specific subset of traffic defined by the **encap-match** command.

The **no** form of this command removes the MACsec instance.

Parameters *port-id* — Specifies the sub-port id index.

Values 1 to 1023

create — Creates a new sub-port.

ca-name

Syntax **ca-name** *ca-name*
no ca-name

Context config>port>ethernet>dot1x>macsec>sub-port

Description This command configures the Connectivity Association (CA) linked to this MACsec sub-port. The specified CA provides the MACsec parameter to be used or negotiated with other peers.

The **no** form of this command removes the CA from the MACsec sub-port.

Parameters *ca-name* — Specifies the appropriate ca to be used under this MACsec sub-port, up to 32 characters.

eapol-destination-address

Syntax **eapol-destination-address** *mac*
no eapol-destination-address

Context config>port>ethernet>dot1x>macsec>sub-port

Description The EAPoL destination MAC address uses a destination multicast MAC address of 01:80:C2:00:00:03. Some networks cannot tunnel this packet over the network and consume these packets, causing the MKA session to fail. This command can change the destination MAC of the EAPoL to the unicast address of the MACsec peer, and as such, the EAPoL and MKA signaling will be unicasted between two peers.

The **no** form of this command returns the value to the default.

Default	no eapol-destination-address
Parameters	<i>mac</i> — Specifies the desired destination MAC address to be used by the EAPOL MKA packets of this sub-port.
Values	aa:bb:cc:dd:ee:ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers.

encap-match

Syntax	encap-match {all-encap double-tag <i>encap-value</i> single-tag <i>encap-value</i> untagged} no encap-match
Context	config>port>ethernet>dot1x>macsec>sub-port
Description	This command defines the sub-set of traffic on this port affected by this MACsec sub-port.

In order to establish an end-to-end communication between the remote MACsec peers encrypting VLAN-tagged traffic, the MKA packets have to be able to travel over the network following the same path as the encrypted traffic. MKA packets are generated with specific tags depending on the traffic match criteria configured, as shown in [Table 49](#).

The **no** form of this command removes all traffic sub-set definitions from the MACsec sub-port.

Table 49 MKA Packet Generation

Configuration	Config Example (<s-tag>.<c-tag>)	MKA Packet Generation	Traffic pattern match/behavior
PORT all-encap	Config>port>ethernet>dot1x>m acsec Sub-port 10 encap-match all-encap ca-name 10	untagged MKA packet	Matches all traffic on the port, including untagged, single-tag, double-tag. This is the Release 15.0 default behavior.
Untagged	Config>port>ethernet>dot1x>m acsec Sub-port 1 encap-match untagged ca-name 2	untagged MKA packet	Matches only untagged traffic on the port

Table 49 MKA Packet Generation (Continued)

Configuration	Config Example (<s-tag>.<c-tag>)	MKA Packet Generation	Traffic pattern match/behavior
802.1Q single S-TAG (specific S-TAG)	Config>port>ethernet>dot1x>m acsec Sub-port 2 encap-match dot1q 1 ca-name 3	MKA packet generated with S-TAG=1	Matches only single-tag traffic on port with tag ID of 1
802.1Q single S-TAG (any S-TAG)	Config>port>ethernet>dot1x>m acsec Sub-port 3 encap-match dot1q * ca-name 4	untagged MKA packet	Matches any single-tag traffic on port
802.1ad double tag (both tag have specific TAGs)	Config>port>ethernet>dot1x>m acsec Sub-port 4 encap-match qinq 1.1 ca-name 5	MKA packet generated with S-tag=1 and C-TAG=1	Matches only double-tag traffic on port with service tag of 1 and customer tag of 1
802.1ad double tag (specific S-TAG, any C-TAG)	Config>port>ethernet>dot1x>m acsec Sub-port 6 encap-match qinq 1.* ca-name 7	MKA packet generated with S-TAG=1	Matches only double-tag traffic on port with service tag of 1 and customer tag of any
802.1ad double tag (any S-TAG, any C-TAG)	Config>port>ethernet>dot1x>m acsec Sub-port 7 encap-match double-tag *.* ca-name 8	untagged MKA packet	Matches any double-tag traffic on port

Default encap-match all-encap

Parameters **all-encap** — Specifies that all traffic patterns are matched including untagged, single-tag or double-tag, and all will be encrypted.

untagged — Specifies that only untagged traffic are matched and encrypted.

single-tag — Specifies that only dot1q traffic are matched. Either all single tag traffic can be matched, by using *, or a specific dot1q tag can be matched.

double-tag — Specifies that only qinq traffic are matched. The service tag can be specifically matched or a wild card match (*.*) can be used.

encap-value — Specifies the type and value of the packet encapsulation to match for this MACsec sub-port.

Type	Parameter
all-encap	—
untagged	—
dot1q	[* s] (s = 0..4094)
qinq	[*.* s.* s.c] (s and c = 0..4094)

- where:
- S = service tag
 - C = customer tag

max-peer

Syntax	max-peer <i>max-peer</i> no max-peer
Context	config>port>ethernet>dot1x>macsec>sub-port
Description	This command configures the max peer allowed under this MACsec instance.



Note: The peer establishment is a race condition and first come first serve. On any security zone, only 32 peers can be supported. See SA Exhaustion Behavior for more details.

The **no** form of this command returns the value to the default.

Default	no max-peer
Parameters	<i>max-peer</i> — The maximum number of peers supported on this port.
Values	0 to 32

shutdown

Syntax	[no] shutdown
Context	config>port>ethernet>dot1x>macsec>sub-port
Description	This command shuts down the MACsec under this sub-port specifically, including MKA negotiation. In the shutdown state, this port is not MACsec capable and all PDUs will be transmitted and expected without encryption and authentication.

The **no** form of this command puts the port in MACsec-enabled mode. A valid CA, different than any other CA configured on any other sub-port of this port and also a *max-peer* value larger than 0 must be configured. In MACsec-enabled mode, packets are sent in clear text until the MKA session is up, and if the **rx-must-be-encrypted** is set on the port, all incoming packets with no MACsec encapsulations are dropped.

Default shutdown

max-auth-req

Syntax **max-auth-req** *max-auth-request*

Context config>port>ethernet>dot1x

Description This command configures the maximum number of times that the router will send an access request RADIUS message to the RADIUS server. If a reply is not received from the RADIUS server after the specified number attempts, the 802.1x authentication procedure is considered to have failed.

The **no** form of this command returns the value to the default.

Default max-auth-req 2

Parameters *max-auth-request* — The maximum number of RADIUS retries.

Values 1 to 10

port-control

Syntax **port-control** [auto | force-auth | force-unauth]

Context config>port>ethernet>dot1x

Description This command configures the 802.1x authentication mode.

The **no** form of this command returns the value to the default.

Default port-control force-auth

Parameters **force-auth** — Disables 802.1x authentication and causes the port to transition to the authorized state without requiring any authentication exchange. The port transmits and receives normal traffic without requiring 802.1x-based host authentication.

force-unauth — Causes the port to remain in the unauthorized state, ignoring all attempts by the hosts to authenticate. The switch cannot provide authentication services to the host through the interface.

auto — Enables 802.1x authentication. The port starts in the unauthorized state, allowing only EAPoL frames to be sent and received through the port. Both the router and the host can initiate an authentication procedure. The port will remain in unauthorized state (no traffic except EAPoL frames is allowed) until the first client is authenticated successfully. After this, traffic is allowed on the port for all connected hosts.

quiet-period

Syntax	quiet-period <i>seconds</i>
Context	config>port>ethernet>dot1x
Description	This command configures the period between two authentication sessions during which no EAPoL frames are sent by the router. The no form of this command returns the value to the default.
Default	quiet-period 60
Parameters	<i>seconds</i> — Specifies the quiet period in seconds. Values 1 to 3600

radius-plcy

Syntax	radius-plcy <i>name</i> no radius-plcy
Context	config>port>ethernet>dot1x
Description	This command configures the RADIUS policy to be used for 802.1x authentication. An 802.1x RADIUS policy must be configured (under config>security>dot1x) before it can be associated to a port. If the RADIUS policy-id does not exist, an error is returned. Only one 802.1x RADIUS policy can be associated with a port at a time. The no form of this command removes the RADIUS policy association.
Default	no radius-plcy
Parameters	<i>name</i> — Specifies an existing 802.1x RADIUS policy name, up to 32 characters.

re-auth-period

Syntax	re-auth-period <i>seconds</i> no re-auth-period

Context	config>port>ethernet>dot1x
Description	This command configures the period after which re-authentication is performed. This value is only relevant if re-authentication is enabled. The no form of this command returns the value to the default.
Default	re-auth-period 3600
Parameters	<i>seconds</i> — Specifies the re-authentication delay period in seconds. Values 1 to 9000

re-authentication

Syntax	[no] re-authentication
Context	config>port>ethernet>dot1x
Description	This command enables/disables periodic 802.1x re-authentication. When re-authentication is enabled, the router will re-authenticate clients on the port every re-auth-period seconds. The no form of the command returns the value to the default.
Default	no re-authentication

server-timeout

Syntax	server-timeout <i>seconds</i> no server-timeout
Context	config>port>ethernet>dot1x
Description	This command configures the period during which the router waits for the RADIUS server to respond to its access request message. When this timer expires, the router will re-send the access request message, up to the specified number times. The no form of this command returns the value to the default.
Default	server-timeout 30
Parameters	<i>seconds</i> — Specifies the server timeout period, in seconds. Values 1 to 300

supplicant-timeout

Syntax	supplicant-timeout <i>seconds</i> no supplicant-timeout
Context	config>port>ethernet>dot1x
Description	This command configures the period during which the router waits for a client to respond to its EAPOL messages. When the supplicant-timeout expires, the 802.1x authentication session is considered to have failed. The no form of this command returns the value to the default.
Default	supplicant-timeout 30
Parameters	<i>seconds</i> — Specifies the server timeout period in seconds. Values 1 to 300

transmit-period

Syntax	transmit-period <i>seconds</i> no transmit-period
Context	config>port>ethernet>dot1x
Description	This command configures the period after which the router sends a new EAPOL request message. The no form of this command returns the value to the default.
Default	transmit-period 30
Parameters	<i>seconds</i> — Specifies the server transmit period in seconds. Values 1 to 3600

tunneling

Syntax	[no] tunneling
Context	config>port>ethernet>dot1x
Description	This command enables the tunneling of untagged 802.1x frames received on a port and is supported only when the dot1x port-control is set to force-auth. 802.1x tunneling is applicable to both Epipe and VPLS services using either a null SAP or a default SAP on a dot1q port. When configured, untagged 802.1x frames will be switched into the service with the corresponding supported SAP.

The **no** form of this command disables tunneling of untagged 802.1x frames.

Default no tunneling

down-when-looped

Syntax **down-when-looped**

Context config>port>ethernet

Description This command configures Ethernet loop detection attributes.

dot1x

Syntax **dot1x**

Context config>port>ethernet

Description This command enables access to the context to configure port-specific 802.1x authentication attributes. This context can only be used when configuring a Fast Ethernet, Gigabit or 10-Gb Ethernet LAN ports on an appropriate MDA.

keep-alive

Syntax **keep-alive** *timer*
no keep-alive

Context config>port>ethernet>dwl

Description This command configures the time interval between keep-alive PDUs.

Default no keep-alive

Parameters *timer* — Specifies the time interval, in seconds, between keep-alive PDUs.

Values 1 to 120

retry-timeout

Syntax **retry-timeout** *timer*
no retry-timeout

Context config>port>ethernet>dwl

Description This command configures the minimum wait time before re-enabling port after loop detection.

Default	no retry-timeout
Parameters	<i>timer</i> — Specifies the minimum wait time before re-enabling port after loop detection.
Values	0, 10 to 160

use-broadcast-address

Syntax	[no] use-broadcast-address
Context	config>port>ethernet>dl
Description	This command specifies whether or not the down when looped destination MAC address is the broadcast address, or the local port MAC address, as specified in the port's MAC address.

2.20.2.16 LLDP Port Commands

lldp

Syntax	lldp
Context	config>port>ethernet
Description	This command enables the context to configure Link Layer Discovery Protocol (LLDP) parameters on the specified port.

dest-mac

Syntax	dest-mac <i>bridge-mac</i>
Context	config>port>ethernet>lldp
Description	This command configures destination MAC address parameters.
Parameters	<i>bridge-mac</i> — Specifies destination bridge MAC type to use by LLDP.
Values	
nearest-bridge	Specifies to use the nearest bridge.
nearest-non-tpmr	Specifies to use the nearest non-Two-Port MAC Relay (TPMR).
nearest-customer	Specifies to use the nearest customer.

admin-status

Syntax	admin-status { rx tx tx-rx disabled }
Context	config>port>ethernet>lldp>dstmac
Description	This command configures LLDP transmission/reception frame handling.
Default	admin-status disabled
Parameters	rx — Specifies the LLDP agent will receive, but will not transmit LLDP frames on this port. tx — Specifies that the LLDP agent will transmit LLDP frames on this port and will not store any information about the remote systems connected. tx-rx — Specifies that the LLDP agent transmits and receives LLDP frames on this port. disabled — Specifies that the LLDP agent does not transmit or receive LLDP frames on this port. If there is remote systems information which is received on this port and stored in other tables, before the port's admin status becomes disabled, then the information will naturally age out.

notification

Syntax	[no] notification
Context	config>port>ethernet>lldp>dstmac
Description	This command enables LLDP notifications. The no form of the command disables LLDP notifications.
Default	no notification

port-id-subtype

Syntax	port-id-subtype { tx-if-alias tx-if-name tx-local }
Context	config>port>ethernet>lldp>dstmac
Description	This command specifies how to encode the PortID TLV transmit to the peer. Some releases of the NSP NFM-P require the PortID value require the default if-Alias in order to properly build the Layer Two topology map using LLDP. Selecting a different option will impact the NSP NFM-P's ability to build those Layer Two topologies.
Default	port-id-subtype tx-local
Parameters	tx-if-alias — Transmits the ifAlias String (subtype 1) that describes the port as stored in the IF-MIB, either user configured or the default entry (i.e. 10/100/Gig Ethernet SFP).

tx-if-name — Transmits the ifName string (subtype 5) that describes the port as stored in the IF-MIB ifName info.

tx-local — The interface ifIndex value (subtype 7) as the PortID.

tunnel-nearest-bridge

Syntax	[no] tunnel-nearest-bridge
Context	config>port>ethernet>lldp>dstmac
Description	The command allows LLDP packets received on the port with the destination address of the nearest bridge to be tunneled without being intercepted on the local port. The dest-mac nearest-bridge must be disable for tunneling to occur. This is applicable to NULL SAP ePipe and VPLS services only.
Default	no tunnel-nearest-bridge

tx-mgmt-address

Syntax	tx-mgmt-address [system] [system-ipv6] [oob] [oob-ipv6] no tx-mgmt-address
Context	config>port>ethernet>lldp>dstmac
Description	This command specifies which management address to transmit. The operator can choose to send the system IPv4 address, the system IPv6 address, the out-of-band IPv4 address, the out-of-band IPv6 address, or any combination of these. The system address is sent only once. The address must be configured for the specific version of the protocol in order to send the management address.
Default	no tx-mgmt-address
Parameters	system — Specifies to use the system IPv4 address. system-ipv6 — Specifies to use the system IPv6 address. oob — Specifies to use the out-of-band IPv4 address for active CPM. oob-ipv6 — Specifies to use the out-of-band IPv6 address for active CPM.

tx-tlvs

Syntax	tx-tlvs [port-desc] [sys-name] [sys-desc] [sys-cap] no tx-tlvs
Context	config>port>ethernet>lldp>dstmac

Description	This command specifies which LLDP TLVs to transmit. The TX TLVs, defined as a bitmap, includes the basic set of LLDP TLVs whose transmission is allowed on the local LLDP agent by the network management. Each bit in the bitmap corresponds to a TLV type associated with a specific optional TLV. Organizationally-specific TLVs are excluded from the this bitmap. There is no bit reserved for the management address TLV type since transmission of management address TLVs are controlled by another object. The no form of the command resets the value to the default.
Default	no tx-tlvs
Parameters	port-desc — Indicates that the LLDP agent should transmit port description TLVs. sys-name — Indicates that the LLDP agent should transmit system name TLVs. sys-desc — Indicates that the LLDP agent should transmit system description TLVs. sys-cap — Indicates that the LLDP agent should transmit system capabilities TLVs.

2.20.2.17 Network Port Commands

network

Syntax	network
Context	config>port>ethernet config>port>sonet-sdh>path config>port>tdm>ds1 config>port>tdm>ds3 config>port>tdm>e1 config>port>tdm>e3
Description	This command enables access to the context to configure network port parameters.

accounting-policy

Syntax	accounting-policy <i>policy-id</i> no accounting-policy
Context	config>port>ethernet config>port>ethernet>access>egr>qgrp config>port>ethernet>access>ing>qgrp config>port>ethernet>network>egr>qgrp config>port>ethernet>network config>port>sonet-sdh>path>network

```
config>port>tdm>ds1>channel-group>network
config>port>tdm>ds3>network
config>port>tdm>e1>channel-group>network
config>port>tdm>e3>network
```

Description This command configures an accounting policy that can apply to an interface.

An accounting policy must be configured before it can be associated to an interface. If the accounting *policy-id* does not exist, an error is returned.

Accounting policies associated with service billing can only be applied to SAPs. Accounting policies associated with network ports can only be associated with interfaces. Only one accounting policy can be associated with an interface at a time.

The **no** form of this command removes the accounting policy association from the network interface, and the accounting policy reverts to the default.

Default No accounting policies are specified by default. You must explicitly specify a policy. If configured, the accounting policy configured as the default is used.

Parameters *policy-id* — The accounting *policy-id* of an existing policy. Accounting policies record either service (access) or network information. A network accounting policy can only be associated with the network port configurations. Accounting policies are configured in the **config>log>accounting-policy** context.

Values 1 to 99

collect-stats

Syntax [no] collect-stats

Context

```
config>port>ethernet>access>egr>qgrp
config>port>ethernet>access>ing>qgrp
config>port>ethernet>network>egr>qgrp
config>port>ethernet>network
config>port>ethernet
config>port>sonet-sdh>path>network
config>port>tdm>ds1>channel-group>network
config>port>tdm>ds3>network
config>port>tdm>e1>channel-group>network
config>port>tdm>e3>network
```

Description This command enables the collection of accounting and statistical data for the network interface. When applying accounting policies, the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated by the XCM/IOM cards, however, the CPU does not obtain the results and write them to the billing file.

If the **collect-stats** command is issued again (enabled), then the counters written to the billing file will include the traffic collected while the **no collect-stats** command was in effect.

Default no collect-stats

queue-policy

Syntax **queue-policy** *name*
no queue-policy

Context config>port>ethernet>network
config>port>sonet-sdh>path>network
config>port>tdm>ds1>channel-group>network
config>port>tdm>ds3>network
config>port>tdm>e1>channel-group>network
config>port>tdm>e3>network

Description This command specifies the existing network queue policy which defines queue parameters such as CBS, high priority only burst size, MBS, CIR and PIR rates, as well as forwarding-class to queue mappings. The network-queue policy is defined in the **config>qos>network-queue** context.

Default queue-policy default

Parameters *name* — Specifies an existing network-queue policy name. The name can be up to 32 characters.

2.20.2.18 HSQ Port Commands

hs-turbo

Syntax	[no] hs-turbo
Context	config>port>ethernet>access>egress>queue-group config>port>ethernet>network>egress>queue-group
Description	This command enables HS turbo queues which allows the corresponding HSQ queue group queues to achieve a higher throughput. The hs-turbo command is not applicable to 10G ports and is ignored when configured under a queue group instance on a 10G port. This command is only applicable to the HSQ IOM (iom4-e-hs) and will fail if configured on all other card types. The no form of the command disables the command.

hs-port-pool-policy

Syntax	hs-port-pool-policy <i>policy-name</i> no hs-port-pool-policy
Context	config>port>ethernet>egress
Description	This command specifies an HS port pool policy to associate with the port egress. An HS port buffer pool policy defines and sizes the port-class buffer pools on an HSQ IOM egress port. A single HS port pool policy is supported per port egress. This command is only applicable to the HSQ IOM (iom4-e-hs) and will fail if configured on all other card types. The no form of the command removes the policy and reapplies the default policy.
Default	hs-port-pool-policy default
Parameters	<i>policy-name</i> — Specifies the HS port pool policy up to 32 characters.

hs-scheduler-overrides

Syntax	hs-scheduler-overrides [create] no hs-scheduler-overrides
Context	config>port>ethernet>egress

Description	This command enables the context to configure HS scheduler overrides which override parameters in the applied HS scheduler policy. This command is only applicable to the HSQ IOM (iom4-e-hs) and will fail if configured on all other card types.
Parameters	create — Keyword used to create HS scheduler overrides. This keyword is requirement and can be enabled or disabled in the environment>create context.

group

Syntax	group <i>group-id</i> rate <i>rate</i> no group <i>group-id</i>
Context	config>port>ethernet>egress>hs-sched-ovr
Description	This command overrides a group rate configured in the HS scheduler policy applied to the port egress. The no form of the command removes the rate override from the port egress configuration.
Parameters	<i>group-id</i> — Specifies the group ID. Values 1 <i>rate</i> — Specifies the maximum rate in megabits per second. When the max keyword follows the rate keyword, the bandwidth limitation is removed from the group. The max keyword is mutually exclusive to the rate parameter. Either the max keyword or a rate value must follow the rate keyword. Values 1 to 100000, max

max-rate

Syntax	max-rate <i>rate</i> no max-rate
Context	config>port>ethernet>egress>hs-sched-ovr
Description	This command overrides the max-rate configured in the HS scheduler policy applied to the port egress. The no form of the command removes the max-rate override from the port egress configuration.
Parameters	<i>rate</i> — Specifies the explicit maximum frame based bandwidth limit, in megabits per second. This parameter is required when executing this command.. Values 1 to 100000, max

scheduling-class

Syntax	scheduling-class <i>class</i> rate <i>rate</i> scheduling-class <i>class</i> weight <i>weight-in-group</i> no scheduling-class <i>class</i>
Context	config>port>ethernet>egress>hs-sched-ovr
Description	This command overrides the scheduling class configuration in the HS scheduler policy applied to the port egress. The scheduling class rate or weight within the WRR group can be overridden. The no form of the command removes the scheduling class override parameters from the port egress configuration.
Parameters	<i>class</i> — Specifies the scheduling class. Values 1 to 6 <i>rate</i> — Specifies the explicit maximum frame based bandwidth limit, in megabits per second, for this HS scheduler policy scheduling class. The rate keyword must be followed by either the keyword max or a rate specified in megabits per second. Values 1 to 100000, max The max keyword specifies that a limit is not enforced for the specified class. The max keyword is mutually exclusive to the rate value and when specified, must directly follow the rate keyword. Setting the rate of the class will fail when the class is currently a member of a group. <i>weight-in-group</i> — Specifies the weight the HS scheduler policy should apply to this scheduling class within the group in which it belongs. The <i>weight-in-group</i> parameter must follow the weight keyword and is used to specify the relative weight of class to the other scheduling classes within the group. Setting the weight will fail if the scheduling class is not currently configured in a group. Values 1 to 127

hs-scheduler-policy

Syntax	hs-scheduler-policy <i>policy-name</i> no hs-scheduler-policy
Context	config>port>ethernet>egress
Description	This command specifies an HS scheduler policy to associate with the port egress which provisions the scheduling behavior of the HSQ scheduler classes. A single HS scheduler policy is supported per port egress. This command is only applicable to the HSQ IOM (iom4-e-hs) and will fail if configured on all other card types.

The **no** form of the command removes the policy and reapplies the default policy.

Default `hs-scheduler-policy default`

Parameters *policy-name* — Specifies the policy name up to 32 characters.

hs-secondary-shaper

Syntax **hs-secondary-shaper** *secondary-shaper-name* [**create**]
no hs-secondary-shaper *secondary-shaper-name*

Context `config>port>ethernet>egress`

Description This command specifies an HS secondary shaper on the port egress. HS secondary shapers are used to apply an aggregate rate and per-scheduling class rates to the set of SAP egress HSQ queue groups which reference them using the SAP egress queue-override **hs-secondary-shaper** command.

By default, the **hs-secondary-shaper** default is applied to each port egress on all HSQ ports and the settings under it can be modified.

Multiple HS secondary shapers are supported per port egress, up to the number supported per-HSQ FP, which is 4096 HS secondary shapers. The number of HS secondary shapers allocated on an HSQ FP can be seen using the **tools dump resource-usage card slot-number fp fp-number** command.

Non-default HS secondary shapers are only configurable on access or hybrid mode ports.

This command is only applicable to the HSQ IOM (iom4-e-hs) and will fail if configured on all other card types.

The **no** form of the command removes the HS secondary shaper from the port egress configuration. An HS scheduler policy cannot be removed when HS scheduler overrides exist on the port egress.

Default `hs-secondary-shaper default`

Parameters *secondary-shaper-name* — Specifies the secondary shaper name up to 32 characters.

aggregate

Syntax [**no**] **aggregate**

Context `config>port>ethernet>egress>hs-sec-shaper`

Description This command enables the context to configure aggregate parameters.

The **no** form of the command removes all of the aggregate parameter values from the configuration of this HS secondary shaper.

low-burst-max-class

Syntax	low-burst-max-class <i>class</i> no low-burst-max-class
Context	config>port>ethernet>egress>hs-sec-shaper>agg
Description	This command specifies which scheduling classes map to the low burst-limit threshold of an egress HS secondary shaper. Egress SAPs can be configured to use an HS secondary shaper that manages their maximum burst limit over a specified aggregate shaping rate. Each HS secondary shaper supports two thresholds, a low burst limit threshold and a high burst limit threshold. By default, all scheduling classes are mapped to the low burst limit threshold. It is important to note that when mapping scheduling classes to the high burst limit threshold an adequate value for the card>fp>egress>hs-fixed-high-thresh-delta must be specified. This is due to the fact that the queues associated with the lower classes may burst over the lower threshold in normal operation due to the scheduler forwarding whole packets. The hs-fixed-high-thresh-delta value should be set to at least two times the maximum frame size to prevent lower threshold class forwarding from also affecting the higher threshold classes when forwarding larger packet sizes. An insufficient high threshold delta defeats the intended purpose of mapping classes to the higher threshold. The system utilizes the lowest value attainable for each low threshold aggregate burst limit without causing shaper underrun conditions. The high burst limit threshold is determined by adding the hs-fixed-high-thresh-delta value configured in the config>card>fp>egress context to the aggregate's low burst limit threshold value. The low-burst-max-class value can be changed at any time for an HS secondary shaper. The no form of the command restores the HS secondary shaper's aggregate low burst limit threshold maximum scheduling class mapping to the default value. This causes all sets of queues associated with the hs-secondary-shaper secondary-shaper-name to have all scheduling classes mapped to the low burst limit threshold.
Default	low-burst-max-class 6
Parameters	<i>class</i> — Specifies the low burst maximum class. This parameter is required when executing the low-burst-max-class command. The parameter reflects the highest scheduling class that will be associated with the low burst limit threshold associated with the HS secondary aggregate shaper. Scheduling classes higher than scheduling class ID will be associated with the high burst limit threshold. Values 1 to 6

rate

Syntax	rate <i>rate</i> no rate

Context	config>port>ethernet>egress>hs-sec-shaper>agg config>port>ethernet>egress>hs-sec-shaper>class
Description	This command specifies the rate allowed for the HS secondary shaper's aggregate rate and per-class rates. The no form of the command reverts to the default.
Default	rate max
Parameters	<i>rate</i> — Specifies the maximum rate in kilobits per second. When the max keyword follows the rate keyword, the bandwidth limitation is removed from the aggregate or class. The max keyword is mutually exclusive to the rate parameter. Either max or a rate value must follow the rate keyword. Values 1 to 100000000, max

class

Syntax	[no] class <i>class-number</i>
Context	config>port>ethernet>egress>hs-sec-shaper
Description	This command specifies the HS secondary shaper class. The no form of the command reverts the classes rate to its default value.
Parameters	<i>class-number</i> — Specifies the HS secondary shaper class identifier. Values 1 to 6

2.20.2.19 Interface Group Handler Commands

interface-group-handler

Syntax	[no] interface-group-handler <i>index</i>
Context	config
Description	This command creates an interface group handler that can be associated with a number of independent IP links. The purpose of the group is to operationally disable all interfaces in a common group if the number of active links drops below the minimum interface threshold. The no form of this command deletes the interface group handler. All members must be removed before the IGH can be deleted.

Parameters *index* — Identifies the specific Interface Group Handler.
Values 1 to 100

member

Syntax **[no] member** *port-id*

Context config>interface-group-handler

Description This command binds the specified port with the associate Interface Group Handler. Up to eight **member** commands can be issued to add multiple ports to the associated IGH. The **member** must be a port or channel on a SONET or POS MDA. It must be a physical port or channel in network mode, and not bound to any router interfaces. A port or channel cannot be a member of more than one IGH at the same time. MLPPP bundles and their members cannot be IGH members.

The **no** form of this command removes the specified port ID from the associated IGH.

Parameters *port-id* — Identifies the port to be associated with the interface group handler.

threshold

Syntax **threshold** *num-members*
no threshold

Context config>interface-group-handler

Description This command identifies the minimum number of active links that must be present for the interface group handler to be active. A threshold of 1 effectively disables the effect of the interface group handler.

The **no** form of this command resets the threshold to 1.

Note that for APS configurations, if the ber-sd or ber-sf threshold rates must be modified, the changes must be performed at the line level on both the working and protect APS port member.

Default threshold 1

Parameters *num-members* — Specifies the minimum number of active links that must be present for the interface group handler to be active.
Values 1 to 8

2.20.2.20 Multilink-Bundle Port Commands

The following Multilink-Bundle Port commands are supported on the 7750 SR only.

multilink-bundle

Syntax	[no] multilink-bundle
Context	config>port
Description	This command creates the context to configure bundle properties for this bundle port.

fragment-threshold

Syntax	fragment-threshold <i>fragment-threshold</i> fragment-threshold unlimited no fragment-threshold
Context	config>port>multilink-bundle
Description	This command sets the maximum length in bytes of a fragment transmitted across a multilink bundle. The no form of this command resets the fragment threshold back to the default value.
Default	fragment-threshold 128
Parameters	<i>fragment-threshold</i> — Specify the maximum fragment length, in bytes, to be transmitted across a multilink bundle. Note that the value range is dependent on the MDA type. For example: channelized MDAs, such as the m1-choc12-sfp, m4-choc3-sfp, m12-chds3, and m4-chds3, support values of 128, 256, 512; ASAP channelized MDAs support any value in the valid range. Values 128 to 512 bytes inclusive for MLPPP and MLFR bundles 128 bytes for IMA bundles unlimited — This keyword disables fragmentation (MLPPP and MLFR only).

interleave-fragments

Syntax	[no] interleave-fragments
Context	config>port>multilink-bundle
Description	This command enables Link Fragmentation and Interleaving on the multilink bundle.

The **no** form of this command disables Link Fragmentation and Interleaving on the multilink bundle.

member

Syntax	[no] member <i>port-id</i>		
Context	config>port>multilink-bundle		
Description	This command binds a channel group to a multilink bundle. For IMA and MLFR groups, this command binds a channel group filling up the entire DS-1 or E-1. For MLPPP groups, fractional (n x ds0) DS1 or E1 links are also allowed. However, fractional DS1 links and fractional E1 links may not be combined in the same multilink bundle. If a channel with a different number of timeslots than the primary-link member is added to the bundle, a warning will be provided. The no form of this command removes the specified channel group from the multilink bundle.		
Parameters	<i>port-id</i> — Specifies the physical port ID in the following format:		
	<i>port-id</i>	<i>slot/mda/port.channel</i>	
	eth-sat-id	esat-id/slot/port	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-id.sub-port	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

minimum-links

Syntax	minimum-links <i>minimum-links</i> no minimum-links
Context	config>port>multilink-bundle
Description	This command sets the minimum number of links that must be active for the bundle to be active. If the number of active links drops below the configured minimum then the multilink bundle will transition to an operationally down state. The no form of this command removes the minimum link limit.
Default	minimum-links 1

Parameters	<i>minimum-link</i> — Specifies the minimum link limit, expressed as an integer.
Values	1 to 8

mlfr

Syntax	mlfr
Context	config>port>multilink-bundle
Description	This command enables the context to configure a Multi-link Frame Relay (MLFR) bundle.

identifier

Syntax	identifier <i>frf16-identifier</i> no identifier
Context	config>port>ml-bundle>mlfr
Description	This command defines the identifier for the MLFR bundle. The no form of this command resets the value to null.
Default	null
Parameters	<i>frf16-identifier</i> — Specifies the bundle ID string. The string can be up to 49 characters long.

identifier

Syntax	identifier <i>frf16-identifier</i> no identifier
Context	config>port>tdm>ds1>channel-group>frame-relay
Description	This command defines the identifier for a frame-relay link when used in an MLFR bundle. The no form of this command resets the value to null.
Default	null
Parameters	<i>frf16-identifier</i> — Specifies the bundle ID string. The string can be up to 49 characters long.

ingress

Syntax	ingress

Context	config>port>ml-bundle>mlfr
Description	This command enables the context to configure the ingress QoS profile for the MLFR bundle.

egress

Syntax	egress
Context	config>port>ml-bundle>mlfr config>port>tdm>ds1>channel-group>frame-relay>frf-12 config>port>tdm>e1>channel-group>frame-relay>frf-12 config>port>tdm>ds3>frame-relay>frf-12 config>port>tdm>e3>frame-relay>frf-12 config>port>sonet-sdh>path>frame-relay>frf-12
Description	This command enables the context to configure the egress QoS profile for an MLFR bundle or a Frame Relay port with FRF.12 UNI/NNI fragmentation enabled.

qos-profile

Syntax	qos-profile <i>profile-id</i> no qos-profile
Context	config>port>ml-bundle>mlfr>ingress config>port>ml-bundle>mlfr>egress config>port>tdm>channel-group>frame-relay>egress config>port>sonet-sdh>path>frame-relay>egress
Description	This command specifies the ingress or egress QoS profile to be used for the configuration of the ingress or egress QoS parameters of an MLFR bundle or a Frame Relay port with FRF.12 UNI/NNI fragmentation enabled. Note that qos-profile on ingress is only applicable to MLFR. The no form of the command removes the parameters from the configuration.
Parameters	<i>profile-id</i> — Specifies the profile number. The value can only be modified if the MLFR bundle or FR port is shut down. Values 1 to 128

frame-relay

Syntax	frame-relay
Context	config>port>ml-bundle>mlfr
Description	This command configures the Frame Relay parameters.

lmi-type

Syntax	lmi-type {ansi itu none rev1}
Context	config>port>multi-link-bundle>mlfr>frame-relay
Description	This command configures the LMI type.
Parameters	ansi — Use ANSI T1.617 Annex D. itu — Use ITU-T Q933 Annex A. none — Disable Frame Relay LMI on the given bundle. rev1 — Use the Rev 1 version of ANSI T1.617 Annex D.

mode

Syntax	mode {dce dte bidir}
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DCE/DTE mode of the Frame Relay interface.
Parameters	dce — Enables the DCE mode. dte — Enables the DTE mode. bidir — Enables the bidirectional mode for LMI types ANSI and ITU.

n391dte

Syntax	n391dte <i>intervals</i> no n391dte
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the number of DTE full status polling intervals for the LMI.
Parameters	<i>intervals</i> — The number of exchanges to be done before requesting a full-status report. A value of 1 specifies to receive full-status messages only. Values 1 to 255

n392dce

Syntax	n392dce <i>threshold</i> no n392dce

Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DCE error threshold for the LMI.
Default	n392dce 3
Parameters	<i>threshold</i> — Specifies the number of errors that will place the bundle in an operationally down state.
Values	1 to 10

n392dte

Syntax	n392dte <i>threshold</i> no n392dte
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DTE error threshold for the LMI.
Parameters	<i>count</i> — Specifies the number of errors that will place the bundle in an operationally down state.
Values	1 to 10

n393dce

Syntax	n393dce <i>count</i> no n393dce
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DCE monitored event count for the LMI.
Parameters	<i>count</i> — Specifies the diagnostic window used to verify link integrity on the DCE interface.
Values	1 to 10

n393dte

Syntax	n393dte <i>count</i> no n393dte
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DTE monitored event count for the LMI.

Parameters	<i>count</i> — Specifies the diagnostic window used to verify link integrity on the DTE interface.
Values	1 to 10

t391dte

Syntax	t391dte <i>keepalive</i> no t391dte
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DTE keepalive timer value for the LMI.
Parameters	<i>keepalive</i> — Specifies the interval in seconds between status inquiries issued by the DTE.
Values	5 to 30

t392dce

Syntax	t392dce <i>keepalive</i> no t392dce
Context	config>port>ml-bundle>mlfr>frame-relay
Description	This command configures the DCE keepalive timer value for the LMI.
Parameters	<i>keepalive</i> — Specifies the expected interval in seconds between status inquiries issued by the DTE equipment.
Values	5 to 30

hello-interval

Syntax	hello-interval <i>timer</i> no hello-interval
Context	config>port>multilink-bundle>mlfr
Description	This command specifies the value of the MLFR bundle T_HELLO timer. The timer controls the rate that hello messages are sent. Following a period of T_HELLO duration, a HELLO message is transmitted onto the bundle link. Note that T_HELLO timer is also used during the bundle link add process as an additional delay before resending an ADD_LINK message to the peer bundle link when the peer bundle link does not answer as expected.

Default	hello-interval 10
Parameters	<i>timer</i> — Specifies the amount of time between HELLO messages in seconds.
Values	1 to 180

ack-timeout

Syntax	ack-timeout <i>timer</i> no ack-timeout
Context	config>port>ml-bundle>mlfr
Description	This command specifies the value of the MLFR bundle T_ACK timer. This timer defines the maximum period to wait for a response to any message sent onto the bundle link before attempting to retransmit a message onto the bundle link.
Default	ack-timeout 4
Parameters	<i>timer</i> — Specifies the wait period in seconds.
Values	1 to 10

retry-limit

Syntax	retry-limit <i>count</i> no retry-limit
Context	config>port>ml-bundle>mlfr
Description	This command specifies the value of the MLFR bundle N_RETRY counter. The counter specifies the number of times a retransmission onto a bundle link will be attempted before an error is declared and the appropriate action taken.
Default	retry-limit 2
Parameters	<i>count</i> — Specifies the number of retransmission attempts.
Values	1 to 5

frf.12

Syntax	frf.12
Context	config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>e1>channel-group>frame-relay

	<pre>config>port>tdm>ds3>frame-relay config>port>tdm>e3>frame-relay config>port>sonet-sdh>path>frame-relay</pre>
Description	This command defines the context to configure the parameters of FRF.12 frame relay fragmentation.

fragment-threshold

Syntax	fragment-threshold <i>fragment-threshold</i> no fragment-threshold
Context	<pre>config>port>tdm>ds1>channel-group>frame-relay>frf.12 config>port>tdm>e1>channel-group>frame-relay>frf.12 config>port>tdm>ds3>frame-relay>frf.12 config>port>tdm>e3>frame-relay>frf.12 config>port>sonet-sdh>path>frame-relay>frf.12</pre>
Description	This command sets the maximum length in bytes of a fragment transmitted across a frame relay port with the FRF.12 UNI/NNI fragmentation enabled. The no form of this command resets the fragment threshold back to the default value.
Default	fragment-threshold 128
Parameters	<i>fragment-threshold</i> — Specifies the maximum fragment length, in bytes, to be transmitted across the FRF.12 port.
Values	128 to 512 bytes

mlppp

Syntax	mlppp
Context	config>port>ml-bundle
Description	This command enables the context to configure multi-link PPP bundle attributes.

egress

Syntax	egress
Context	config>port>ml-bundle>mlppp
Description	This command enables the context to configure egress MLPPP QoS profile parameters for the multilink bundle.

ingress

Syntax	ingress
Context	config>port>ml-bundle>mlppp
Description	This command enables the context to configure ingress MLPPP QoS profile parameters for the multilink bundle.

qos-profile

Syntax	qos-profile <i>profile-id</i> no qos-profile
Context	config>port>ml-bundle>mlppp>egress
Description	This command specifies the egress QoS profile to be used for the outgoing traffic over this MLPPP bundle. The no form of the command removes the parameters from the configuration.
Parameters	<i>profile-id</i> — Specifies the egress QoS profile to be used for the outgoing traffic over this MLPPP bundle. The value can only be modified if the bundle is shutdown. Values 1 to 65535

endpoint-discriminator

Syntax	endpoint-discriminator class { ip-address global-mac-address null } [discriminator-id <i>discriminator-id</i>] no endpoint-discriminator
Context	config>port>ml-bundle>mlppp
Description	This command configures the endpoint-discriminator class and ID. The port must be shutdown to modify command parameters. The no form of the command removes the parameters from the configuration.
Parameters	class — Specifies the Link Control Protocol endpoint discriminator class field type. Values ip-address, global-mac-address, null Default Bundle type Default Physical MLPPP bundle ieee802dot1GlobalMacAddress MLPPP bundle protection group IP address

discriminator-id — Specifies the endpoint discriminator identifier value within the specified endpoint-discriminator class. The values below are only for the ip-address class. the string can be up to 20 characters.

Values Any valid IP address.

qos-profile

Syntax	qos-profile <i>profile-id</i> no qos-profile
Context	config>port>ml-bundle>mlppp>ingress
Description	This command specifies the ingress QoS profile to be used for the incoming traffic over this MLPPP bundle.
Parameters	<i>profile-id</i> — Specifies the ingress QoS profile to be used for the incoming traffic over this MLPPP bundle. The value can only be modified if the bundle is shutdown. Values 1 to 65535

magic-number

Syntax	[no] magic-number
Context	config>port>ml-bundle>mlppp
Description	This command allows loopback detection to be enabled and disabled for MLPPP bundles. It is disabled by default. When the magic number option is disabled, the magic number option will not be requested when a member is trying to bring up the LCP layer on a member link; if the remote peer requests this option, it will be rejected. When transmitting echo-requests a magic number of 0 is used. When responding to echo-requests a magic number of 0 is sent. The magic number option is sent to the remote peer during protocol negotiation. If this option is rejected by the remote peer, the router will bring the link up but will be unable to detect loopbacks since the router will always send a magic number of 0 in the echo messages. If this option is accepted by the remote peer, the router will send echo messages with randomly generated magic-numbers. If the SR receives a config-req with the same magic number that was sent out, the router will calculate a new magic number to use and send out another config-request. If the router is persistently seeing the randomly generated magic number in the received config-req, the router will declare a loopback. The no form of the command disables the loopback detection.
Default	no magic-number

multiclass

Syntax	multiclass <i>count</i> no multiclass
Context	config>port>ml-bundle>multiclass
Description	This command enables multi-class MLPPP as defined by RFC 2686, <i>The Multi-Class Extension to Multi-Link PPP</i>, on a MLPPP bundle (including MLPPP bundle protection groups) with 2, 3 or 4 classes. For multiclass MLPPP bundles with a non-zero count, the class index takes valid values from 0 to one less than the maximum number of classes inclusive. For example a 4-class MLPPP bundle has 4 classes with indexes 0, 1, 2, and 3. A bundle must be shutdown with no links for this value to be changed. Entries are created and deleted by the system depending on the number of classes being used by a given MLPPP bundle. The no form of the command disables multi-class MLPPP.
Default	multiclass 4
Parameters	<i>count</i> — Specifies the number of classes in a MLPPP bundle. Values 2 to 4

stateless-aps-switchover

Syntax	[no] stateless-aps-switchover
Context	config>port>ml-bundle> mlppp
Description	This command specifies whether the bundle will perform a stateful or a stateless APS switchover. The value can be changed for APS bundle protection groups of type MLPPP. A stateless switchover implies that PPP is re-negotiated on each member link after the switchover. PPP negotiations may take a few seconds to complete. A stateful switchover implies that after an APS switchover the PPP state of the bundle will be restored based on the bpgpr bundle state before the switchover. The state cannot be changed for normal MLPPP bundles (only applicable for bpgprs). The no form of the command disables stateless APS switchover.
Default	no stateless-aps-switchover

mrru

Syntax	mrru <i>mrru</i> no mrru
Context	config>port>multilink-bundle
Description	This command specifies the maximum received reconstructed unit (MRRU), similar to a maximum transmission unit (MTU), but applies only to MLPPP multilink bundles. The MRRU is the maximum frame size that can be reconstructed from multilink fragments. This command is only valid for MLPPP bundles. The no form of this command resets the MRRU to the default.
Default	mrru 1524
Parameters	<i>mrru</i> — Specifies the maximum received reconstructed unit size, in bytes.
Values	1500 to 9206

protect-bundle

Syntax	[no] protect-bundle <i>bundle-id</i>
Context	config>port>multilink-bundle
Description	This command configures a protect bundle that is part of this BPGrp.
Parameters	<i>bundle-id</i> — Specifies the protection multilink bundle in the bundle protection group. The command syntax must be configured as follows: Values <i>bundle-type-slot/mda.bundle-num</i> bundle-PPP or IMA-slot/mda.bundle-num Creates an MLPPP or IMA bundle. where: bundle: keyword <i>slot:</i> IOM/MDA slot numbers <i>bundle-num:</i> 1 to 336

For example:
router1>config>port>ml-bundle> protect-bundle bundle-ima-1/1.1

red-differential-delay

Syntax	red-differential-delay <i>red-diff-delay</i> [down] no red-differential-delay

Context	config>port>multilink-bundle
Description	This command sets the maximum acceptable differential delay for individual links within a multilink bundle. The differential delay is calculated as the round-trip differential delay for MLPPP bundles, and as uni-directional differential delay for IMA bundles. The no form of this command restores the red-differential-delay defaults.
Parameters	<i>red-diff-delay</i> — Specify the maximum red differential delay value. Values 0 to 50 milliseconds for IMA bundles 0 to 25 milliseconds for all other bundles down — Transition the link that exceeded the differential delay to a down state (for example, remove it from the multilink bundle from an operational perspective).

short-sequence

Syntax	[no] short-sequence
Context	config>port>multilink-bundle
Description	This command specifies that the Multi-link Point to Point Protocol (MLPPP) bundle should use short (12 bit) sequence numbers instead of the default 24-bit sequence number. This command is only valid for MLPPP bundles. The no form of this command disables the short-sequence feature.
Default	no short-sequence

working-bundle

Syntax	[no] working-bundle <i>bundle-id</i>
Context	config>port>multilink-bundle
Description	This command configures a working bundle that is part of this BPGp.
Parameters	<i>bundle-id</i> — Specifies the working multilink bundle in the bundle protection group. The command syntax must be configured as follows: Values <i>bundle-type-slot/mda.bundle-num</i> bundle-PPP or IMA-slot/ mda.bundle-num Creates an MLPPP or IMA bundle. where: bundle: keyword <i>slot:</i> IOM/MDA slot numbers <i>bundle-num:</i> 1 to 336

For example:

```
router1>config>port>ml-bundle> working-bundle bundle-ima-1/1.1
```

yellow-differential-delay

Syntax	yellow-differential-delay <i>yellow-diff-delay</i> no yellow-differential-delay
Context	config>port>multilink-bundle
Description	This command sets the yellow warning threshold for the differential delay for members within a multilink bundle. If circuit's delay exceeds the yellow-differential delay value, a log message and SNMP trap is sent. This command is only valid for MLPPP bundles. The differential delay is calculated as the round-trip differential delay for MLPPP bundles. The no form of this command removes the yellow-differential-delay.
Parameters	<i>yellow-diff-delay</i> — Specifies the maximum yellow differential delay threshold value, in milliseconds.
Values	1 to 25

ima

Syntax	ima
Context	config>port>multilink-bundle
Description	This command enables the context to configure parameters for an Inverse Multiplexing over ATM (IMA) group. An IMA group is a collection of physical links bundled together and assigned to an ATM interface. IMA enables a high-speed channel that is composed of ATM cells to be transported as a number of lower-speed circuits. Then they are reassembled as the original high-speed ATM channel. This command is only valid for IMA bundles.

link-delay

Syntax	link-delay { activate deactivate } <i>milli-seconds</i> no link-delay { activate deactivate }
Context	config>port>ml-bundle>ima
Description	This command specifies the time to delay between detection of a link activation/deactivation condition and acting upon it (going in/out of the RX failure state on a link).

Parameters	activate <i>milli-seconds</i> — Specifies the time, in milli-seconds, used to clear an existing LIF or LODS alarm. The time specified determines how long is needed for member links to stabilize before being activated.
	Values 1 to 30000
	Default 10000
	deactivate <i>milli-seconds</i> — Specifies the time, in milli-seconds, used to raise an LIF or LODS alarm. The time specified determines how long before a member link is declared in error and is deactivated.
	Values 1 to 30000
	Default 2000

max-bandwidth

Syntax	max-bandwidth <i>number-links</i> no max-bandwidth
Context	config>port>ml-bundle>ima
Description	This command specifies the number of links used to determine the maximum configurable bandwidth that is allowed to be used for this IMA group. The maximum bandwidth is computed as: Maximum Configurable ATM Bandwidth (MCAB) = $(\text{number-links}) * (M-1)/M * (2048/2049) * \text{primary member link speed}$ Where: - M is the IMA frame size (128) - Primary member link speed is either E-1 — 1920 kb/s or DS-1 — 1539 kb/s. E-1 speed is used for a group with no members. The total ATM bandwidth of services over shaped VCs cannot exceed the MCAB value as result of adding more services or removing member links. The no form of the command resets the max-bandwidth to its default value
Default	max-bandwidth 8
Parameters	<i>number-links</i> — Specifies the number of links that determines the maximum configurable bandwidth that is allowed to be used for this IMA group. Values 1 to 8

test-pattern-procedure

Syntax	test-pattern-procedure
Context	config>port>ml-bundle>ima
Description	This command enables the context to configure IMA test pattern procedures. Note that this command and sub-commands are not saved in the router configuration between reboots.

test-link

Syntax	test-link <i>port-id</i> no test-link
Context	config>port>ml-bundle>ima>test-pattern-procedure
Description	This command specifies IMA members on which an IMA test pattern procedure is to be performed. The no form of this command deletes the link from test-pattern procedure. The test-pattern procedure must be shutdown first.
Default	no test-link
Parameters	<i>port-id</i> — Specifies the port ID that verifies link connectivity within an IMA group.
Values	
<i>port-id</i>	<i>slot/mda/port</i> [<i>channel</i>]
eth-sat-id	<i>esat-id/slot/port</i> <i>esat</i> keyword <i>id</i> 1 to 20
pxc-id	<i>pxc-id.sub-port</i> <i>pxc</i> keyword <i>id</i> 1 to 64 <i>sub-port</i> a, b
aps-id	<i>aps-group-id[.channel]</i> <i>aps</i> keyword <i>group-id</i> 1 to 128

test-pattern

Syntax	test-pattern <i>pattern</i> no test-pattern
Context	config>port>ml-bundle>ima>test-pattern-procedure

Description	This command specifies the transmit test pattern in an IMA group loopback operation. This value can only be changed when the test-pattern-procedure command is shut down The no form of this command restores the test-pattern to the default.
Default	test-pattern 0
Parameters	<i>pattern</i> — Specifies an integer taking the following values: Values 0 to 255

shutdown

Syntax	[no] shutdown
Context	config>port>ml-bundle>ima>test-pattern-procedure
Description	This command enables a configured IMA test pattern procedure. The no form of this command disables the IMA test pattern procedure.

version

Syntax	version <i>IMA-version</i> no version
Context	config>port>ml-bundle>ima
Description	This command configures the IMA version for the multilink bundle group. If there is a version mismatch between this IMA group and the far end IMA group, the IMA group will become operationally down. Automatic version changing is not supported. To change the IMA version, all member links must be removed from the group first.
Default	version 1-1
Parameters	<i>IMA-version</i> — Specifies the IMA version for this group. Values 1-0: IMA version 1-0 1-1: IMA version 1-1

2.20.2.21 SONET/SDH Port Commands

sonet-sdh

Syntax	sonet-sdh

Context config>port

Description This command enables access to the context to configure SONET/SDH ports. This context can only be used when configuring an OC-3, OC-12, OC-48, OC-192, and OC-768 SONET/SDH ports on an appropriate MDA.

This command also enables access to the context to configure SONET/SDH parameters for an Ethernet port in WAN PHY (xgig wan) mode.

The 10 Gigabit Ethernet LAN port also has SONET/SDH characteristics. However, these characteristics are predetermined and not configurable.

This command is supported by TDM satellite.

clock-source

Syntax clock-source {loop-timed | node-timed}

Context config>port>sonet-sdh

Description This command configures the clock to be used for transmission of data out towards the line. The options are to use the locally recovered clock from the line's receive data stream or the node central reference.

When changing the clock source for a port on an OC-48 MDA, a brief transmit interruption can occur on all ports of that MDA. Note that all SONET/SDH MDAs/CMAs support loop timing.

The **node-timed** parameter in this command is supported by TDM satellite.

Parameters **loop-timed** — The link recovers the clock from the received data stream.
node-timed — The link uses the internal clock when transmitting data.

framing

Syntax framing {sonet | sdh}

Context config>port>sonet-sdh

Description This command specifies SONET/SDH framing to be either SONET or SDH.

This command is supported by TDM satellite.

Default framing sonet

Parameters **sonet** — Configures the port for SONET framing.
sdh — Configures the port for SDH framing.

group

Syntax	group <i>sonet-sdh-index</i> payload { tu3 vt2 vt15 }
Context	config>port>sonet-sdh
Description	This command configures payload of the SONET/SDH group. This command is supported by TDM satellite, however the tu3 parameter is not. For example: <pre>config>port>sonet-sdh# group tug3-1.1 payload tu3 group tug3-1.2 payload vt2 group tug3-1.3 payload vt2 group tug3-2.1 payload vt15 group tug3-2.2 payload vt15 group tug3-2.3 payload tu3 group tug3-3.1 payload tu3 group tug3-3.2 payload tu3 group tug3-3.3 payload tu3</pre>
Parameters	<i>sonet-sdh-index</i> — Specifies the components making up the specified SONET/SDH path. Depending on the type of SONET/SDH port the <i>sonet-sdh-index</i> must specify more path indexes to specify the payload location of the path. tu3 — Specifies the Tributary Unit Group (TUG3) on a path. Configures the port or channel for transport network use. vt2 — Configures the path as a virtual tributary group of type vt2. vt15 — Configures the path as a virtual tributary group of type vt15.

hold-time

Syntax	hold-time {[up <i>hold-time-up</i>] [down <i>hold-time-down</i>]} no hold-time
Context	config>port>sonet-sdh
Description	This command configures SONET link dampening timers in 100s of milliseconds. This guards against reporting excessive interface transitions. This is implemented by not advertising subsequent transitions of the interface to upper layer protocols until the configured timer has expired.

Note: For APS configurations, the **hold-time down** and **up** default values are 100 ms and 500 ms respectively. If there is a large communication delay (time to exchange K1/K2 bytes) between the APS Controllers of the two endpoints of an APS link, it is highly suggested to increase the default hold-time down timer on the APS group port accordingly with the communication delay. See the [aps](#) command for more information.

This command is supported by TDM satellite.

Default	no hold-time
Parameters	up <i>hold-time-up</i> — Configures the hold-timer for link up event dampening. A value of zero (0) indicates that an up transition is reported immediately. Values 0 to 100 down <i>hold-time-down</i> — The hold-timer for link down event dampening. A value of zero (0) indicates that a down transition is reported immediately. Values 0 to 100

loopback

Syntax	loopback { line internal } no loopback
Context	config>port>sonet-sdh
Description	This command activates a loopback on the SONET/SDH port. The SONET port must be in a shut down state to activate any type of loopback. The loopback setting is never saved to the generated/saved configuration file. Note that loopback mode changes on a SONET/SDH port can affect traffic on the remaining ports. This command is supported by TDM satellite.
Default	no loopback
Parameters	line — Set the port into line loopback state. internal — Set the port into internal loopback state.

report-alarm

Syntax	[no] report-alarm [loc] [lais] [lrdi] [ss1f] [lb2er-sd] [lb2er-sf] [slof] [slos] [lrei]
Context	config>port>sonet-sdh

Description	This command enables logging of SONET (SDH) line and section alarms for a SONET-SDH port. Only line and section alarms can be configured in the SONET/SDH context, for path alarms see the sonet-sdh>path context. The no form of this command disables logging of the specified alarms. This command is supported on TDM satellites.
Parameters	loc — Reports a loss of clock which causes the operational state of the port to be shut down. Default loc alarms are issued lais — Reports line alarm indication signal errors. When configured, lais alarms are raised and cleared. Default lais alarms are not issued lrdi — Reports line remote defect indication errors. LRDI's are caused by remote LOF, LOC, LOS. When configured, lrdi alarms are raised and cleared. Default lrdi alarms are issued ss1f — Reports section synchronization failure which is detected when the S1 byte is not consistent for 8 consecutive frames. When configured, ss1f alarms are raised and cleared. Default ss1f alarms are not issued lb2er-sd — Reports line signal degradation BER (bit interleaved parity) errors. Use the threshold command to set the error rate(s) that when crossed determine signal degradation and signal failure. When configured, lb2er-sd alarms are raised and cleared. Default lb2er-sd alarms are not issued lb2er-sf — Reports line signal failure BER errors. Use the threshold command to set the error rate(s) that when crossed determine signal degradation and signal failure. When configured, lb2er-sf alarms are raised and cleared. Default lb2er-sf alarms are issued slof — Reports section loss of frame errors. When configured, slof alarms are raised and cleared. Default slof alarms are issued slos — Reports a section loss of signal error on the transmit side. When configured, slos alarms are raised and cleared. Default slos alarms are issued lrei — Reports a line error condition raised by the remote as a result of b1 errors received from this node. When configured, lrei traps are raised but not cleared. Default lrei traps are not issued

reset-port-on-path-down

Syntax	[no] reset-port-on-path-down
Context	config>port>sonet-sdh
Description	This command configures whether the SONET/SDH port will reset when the path transitions to an operationally down state. This command only affects SONET/SDH ports on 7750 4-port OC48 SFP “-B” MDAs.
Default	no reset-port-on-path-down

section-trace

Syntax	section-trace {increment-z0 byte <i>value</i> string <i>string</i>}				
Context	config>port>sonet-sdh				
Description	This command configures the section trace bytes in the SONET section header to inter-operate with some older versions of ADMs or regenerators that require an incrementing STM ID. You can explicitly configure an incrementing STM value rather than a static one in the SDH overhead by specifying the z0-increment. This command is supported on TDM satellite.				
Default	section-trace byte 0x1				
Parameters	<i>increment-z0</i> — Configures an incrementing STM ID instead of a static value. <i>value</i> — Sets values in SONET header bytes. <table><tr><td>Default</td><td>0x1</td></tr><tr><td>Values</td><td>0 to 255 or 0x00 to 0xFF</td></tr></table> <i>string</i> — Specifies a text string that identifies the section. The string can be a maximum of 16 characters.	Default	0x1	Values	0 to 255 or 0x00 to 0xFF
Default	0x1				
Values	0 to 255 or 0x00 to 0xFF				

speed

Syntax	speed {oc3 oc12} no speed
Context	config>port>sonet-sdh
Description	This command configures the speed of a SONET/SDH port as either OC3 or OC12. The framer for this MDA operates in groups of four. Changing the port speed for a port requires resetting the framer and causes a slight disruption on all four ports. The first framer controls ports 1,2,3,4, the second framer controls ports 5,6,7,8 and so on.

To change the port speed on a SONET/SDH port, the port must be administratively shut down and all channels must be removed. When the port speed is changed, the default channel configuration is recreated.

The **no** form of this command reverts back to default.

This command is supported on TDM satellite.

Default	speed oc12
Parameters	oc3 — Sets the speed of the port to OC-3. oc12 — Sets the speed of the port to OC-12.

suppress-lo-alarm

Syntax	[no] suppress-lo-alarm
Context	config>port>sonet-sdh
Description	This command enables the suppression of lower order alarms on SONET/SDH port such as MLPPP bundle alarms, DS1/E1 links alarms and 336 APS channel groups alarms. The no form of the command disables the suppression of lower order alarms on SONET/SDH port.

tx-dus

Syntax	[no] tx-dus
Context	config>port>ethernet>ssm config>port>sonet-sdh
Description	This command forces the QL value transmitted from the SSM channel of the SONET/SDH port or the Synchronous Ethernet port to be set to QL-DUS/QL-DNU. This capability is provided to block the use of the interface from the SR/ESS for timing purposes. This command is supported on TDM satellite.
Default	no tx-dus

threshold

Syntax	threshold {ber-sd ber-sf} rate <i>threshold-rate</i> no threshold {ber-sd ber-sf}
Context	config>port>sonet-sdh

Description	This command configures the line signal degradation bit error rate (BER) and line signal failure thresholds. Line signal (b2) bit interleaved parity error rates are measured and when they cross either the degradation or failure thresholds alarms are raised (see the report-alarm command), furthermore if the failure threshold is crossed the link will be set to operationally down. For APS configurations, if the ber-sd or ber-sf threshold rates must be modified, the changes must be performed at the line level on both the working and protect APS port member. The no form of this command reverts to the default value.
Default	threshold ber-sd rate 6 — Signal degrade BER threshold of 10-6. threshold ber-sf rate 3 — Signal failure BER threshold of 10-3.
Parameters	ber-sd — Specifies the BER that specifies signal degradation. ber-sf — Specifies the BER that specifies signal failure. <i>threshold-rate</i> — The BER negative exponent (n in 10-n), expressed as a decimal integer.
Values	3 to 9 (10-3 to 10-9) for ber-sd, 3 to 6 for ber-sf

2.20.2.22 SONET/SDH Path Commands

path

Syntax	[no] path [sonet-sdh-index]						
Context	config>port>sonet-sdh						
Description	This command defines the SONET/SDH path. The no form of this command removes the specified SONET/SDH path. This command is supported on TDM satellite.						
Default	full channel (or clear channel)						
Parameters	<i>sonet-sdh-index</i> — Specifies the components making up the specified SONET/SDH path. Depending on the type of SONET/SDH port the <i>sonet-sdh-index</i> must specify more path indexes to specify the payload location of the path. The <i>sonet-sdh-index</i> differs for SONET and SDH ports. <table><tr><td>Values</td><td>sts192 (for the 7950 XRS only)</td></tr><tr><td></td><td>sts1-x.x (for the 7450 ESS and 7750 SR), tu3, vt2, vt15</td></tr></table> <table><tr><td>SONET</td><td>SDH</td></tr></table>	Values	sts192 (for the 7950 XRS only)		sts1-x.x (for the 7450 ESS and 7750 SR), tu3, vt2, vt15	SONET	SDH
Values	sts192 (for the 7950 XRS only)						
	sts1-x.x (for the 7450 ESS and 7750 SR), tu3, vt2, vt15						
SONET	SDH						

OC-192	STS-48-index	STM-64	AUG-16-index
	STS-12-index		AUG-4-index
	STS-3-index		AUG-1-index
	STS-1-index		AU-3-index
OC-48	STS-12-index	STM-16	AUG-4-index
	STS-3-index		AUG-1-index
	STS-1-index		AU-3-index
OC-12	STS-3-index	STM-4	AUG-1-index
	STS-1-index		AU-3-index
OC-3	STS-1-index	STM-1	AU-3-index

In addition the support of virtual tributary circuits adds an additional level of complexity and several addition levels of indexes.

payload

Syntax	payload { sts3 tug3 ds3 e3 vt2 vt15 ds1 e1 }
Context	config>port>sonet-sdh>path
Description	This command specifies if the associated SONET/SDH path is an asynchronous circuit or a virtual tributary group (VT). This command is only applicable to channelized MDAs. This command is supported on TDM satellite, however the sts3, ds3, and e3 parameters are not supported.
Parameters	sts3 — Configures STS3/STM1 payload as clear channel. tu3 — Configures STS3/STM1 payload as Tributary Unit Group 3 (TUG3). ds3 — Configures the port or channel as DS-3 STS1/VC3 payload as DS-3. e3 — Configures the port or channel as E-3 STS1/VC3 payload as E-3. vt2 — Configures the path STS1 payload as vt2 as a virtual tributary group. Only allowed on STS-1 nodes (SONET VT container). vt15 — Configures the path as a virtual tributary group. Only allowed on STS-1 nodes (SONET VT container). ds1 — Configures the port or channel as DS1.vt15 or vt2 payload as DS-1. e1 — Configures VT2 payload as E-1.

keepalive

Syntax	keepalive <i>time-interval</i> [dropcount <i>drop-count</i>]

no keepalive

Context	config>port>sonet-sdh>path>ppp
Description	This command configures the time interval at which keepalive requests are issued.
Parameters	<i>time-interval</i> — Specifies the interval used to send periodic keepalive packets. Values 1 to 60 seconds <i>drop-count</i> — Specifies the number of consecutive keepalive failed request attempts or remote replies that can be missed after which the port is operationally downed. Values 1 to 255

report-alarm

Syntax	[no] report-alarm [pais] [plop] [prdi] [pplm] [prei] [puneq] [plcd]
Context	config>port>sonet-sdh>path
Description	This command enables logging of SONET (SDH) path alarms for a SONET-SDH port. Only path alarms can be configured in the channel context. The no form of this command disables logging of the specified alarms. This command is supported on TDM satellites.
Parameters	pais — Reports path alarm indication signal errors. When configured, pais alarms are raised and cleared. Default pais alarms are not issued plop — Reports path loss of pointer (per tributary) errors. When configured, plop traps are raised but not cleared. Default plop traps are issued prdi — Reports path remote defect indication errors. When configured, prdi alarms are raised and cleared. Default prdi alarms are not issued pplm — Reports a path payload mismatch, as a result the channel will be operationally downed. When configured, pplm traps are raised but not cleared. Default pplm traps are issued prei — Reports a path error condition raised by the remote as a result of b3 errors received from this node. When configured, prei traps are raised but not cleared. Default prei traps are not issued puneq — Reports path unequipped errors. Reports path unequipped signal errors. Default puneq traps are issued

plcd — Reports path loss of code group delineation errors. It is applicable only when the value of **xgig** is set to WAN.

Default **plcd** traps are not issued

crc

Syntax **crc {16 | 32}**

Context config>port>sonet-sdh>path

Description A 16 bit CRC can only be configured on an OC-3 channel, all other channel speeds must use a 32 bit CRC except for the paths configured with **encap-type atm** at OC3 speed.

Default crc 16 for OC-3, DS-1, DS-3
crc 32 for OC-12, OC-48, ATM-OC12/3, AT-MOC-3, and so on



Note: The CRC default is 32 when the **encap-type** is set to ATM and also, the default cannot be changed when the **encap-type** is set to ATM.

Parameters **16** — Use 16 bit checksum for the associated port/channel.

32 — Use 32 bit checksum for the associated port/channel.

encap-type

Syntax **encap-type {atm | bcp-null | bcp-dot1q | ipcp | ppp-auto | frame-relay | wan-mirror | cisco-hdlc}**

Context config>port>sonet-sdh>path

Description This command configures the encapsulation method used to distinguish customer traffic on an access SONET/SDH channel sub-port.

When the **encap-type** is set to ATM the CRC default cannot be changed.

When the **encap-type** is ATM, ATM sub-layer verification (GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs)*) is automatically enabled. The result of the verification includes:

- Out of Cell Delineation (OCD) event count. The OCD event count is described in RFC 2515, *Definitions of Managed Objects for ATM Management*. Note that multiple events occurring within a second will be counted as 1 event for ATM and ASAP MDAs as a result of a hardware limit.

- Loss of Cell Delineation defect/alarm. The LCD defect/alarm is defined in RFC 2515, *Definitions of Managed Objects for ATM Management*. When a path is in an LCD defect state, the path's operational status will be down. When a path exits the LCD state, the path's operational status will change to up (assuming nothing else causes the path to stay down). A trap is raised to indicate the LCD status change. Also a P-RDI is sent to indicate the defect to the remote end.

The **encap-type** is only required when configuring a SONET/SDH path for access mode.

The **no** form of this command restores the default.

Default	encap-type bcp-null
Parameters	atm — Specifies that the encapsulation on the port is ATM. bcp-null — Only a single service is configured on this channel and IEEE 802.1Q tags are not used as a service delimiter. Any IEEE 802.1Q tags encountered are regarded as part of the customer payload and transparently forwarded. When bcp-null encapsulation is specified, the PPP Bridge Control Protocol (BCP) is activated and all packets on this access port will be encapsulated in accordance with the BCP protocol. Note that null ports will accept q-tagged frames. bcp-dot1q — Ingress frames carry IEEE 802.1Q tags and the tags are used as service delimiter. Any untagged packets are silently discarded with exception of protocol specific packets. When bcp-dot1q encapsulation is specified, the PPP Bridge Control Protocol (BCP) is activated and all packets on this access port will be encapsulated in accordance with the BCP protocol. ipcp — Ingress frames are encapsulated according to the IP Control Protocol. When ipcp encapsulation is specified, the PPP IP Control Protocol will be activated and only packets that comply with IPCP encapsulation are processed; others are silently discarded. ppp-auto — Enables PPP on the associated port/channel. The activation of ipcp and mpplscp is automatically enabled depending on the protocol configuration. This encap type is only valid on ports/channels in network mode. frame-relay — Enables frame relay on the associated port/channel. wan-mirror — The port is used for mirroring of frame-relay and POS ports. On these ports, no link management protocol would run. cisco-hdlc — Monitors line status on a serial interface by exchanging keepalive request messages with peer network devices.

ppp

Syntax	ppp
Context	config>port>sonet-sdh>path

Description This command enables access to the context to configure the LCP operational parameters for a SONET/SDH Point-to-Point Protocol (PPP) link.

report-alarm

Syntax [no] report-alarm {pais | plop | prdi | pplm | prei}

Context config>port>sonet-sdh>path

Description This command enables logging of SONET (SDH) path alarms for a SONET-SDH port. Only path alarms can be configured in the channel context.

The **no** form of this command disables logging of the specified alarms.

Parameters

pais — Reports path alarm indication signal errors. When configured, **pais** alarms are raised and cleared.

Default pais alarms are not issued

plop — Reports path loss of pointer (per tributary) errors. When configured, **plop** traps are raised but not cleared.

Default plop traps are issued

prdi — Reports path remote defect indication errors. When configured, **prdi** alarms are raised and cleared.

Default prdi alarms are not issued

pplm — Reports a path payload mismatch, as a result the channel will be brought down. When configured, **pplm** traps are raised but not cleared.

Default pplm traps are issued

prei — Reports a path error condition raised by the remote as a result of b3 errors received from this node. When configured, **prei** traps are raised but not cleared.

Default prei traps are not issued

scramble

Syntax [no] scramble

Context config>port>sonet-sdh>path

Description This command enables SONET/SDH payload scrambling. Scrambling randomizes the pattern of 1s and 0s carried in a SONET frame. Rearranging or scrambling the pattern prevents continuous strings of all 1s or all 0s and meets the needs of physical layer protocols that rely on sufficient transitions between 1s and 0s to maintain clocking.

For ATM, this command enables or disables ATM cell-level payload scrambling/descrambling using $x^{43}+1$ polynomial as defined in ITU-T I.432.1. Scrambling is enabled by default for the ATM path/channel. Note that this scrambling is done in addition to SONET/SDH frame scrambling/descrambling, which is always enabled in the framer.

The **no** form of this command disables scrambling.

Default no scramble

signal-label

Syntax **signal-label** *value*
no signal-label

Context config>port>sonet-sdh>path

Description This command sets the C2 byte value. The purpose of this byte is to communicate the payload type being encapsulated by SONET framing.

This command is supported on TDM satellite.

Default signal-label 0xcf

Parameters *value* — Specifies the C2 byte value, expressed as a decimal integer or a value in hex format.

Values 1 to 254 or 0x01 to 0xfe

trace-string

Syntax **trace-string** [*trace-string*]
no trace-string

Context config>port> sonet-sdh>path

Description This command specifies that a J1-path-trace that identifies the circuit is inserted continuously at source. This can be checked against the expected value by the receiver. If no trace string is entered then a null string is used.

The **no** form of this command resets the string to its default.

This command is supported on TDM satellite.

Default The default J1 value is *Alcatel-Lucent* XXX YYY where XXX is the platform number, such as "7750" or "7450", and YYY is the platform acronym, such as "SR" or "ESS". The value does not change when the encap-type changes. The J1 string contains all zeros for a non-provisioned path.

Parameters *trace-string* — Specifies either a string up to 62 bytes for SONET or 15 bytes for SDH. If the string contains spaces, enclose it in quotation marks. String 'zeros' will send all zeros in the J1 bytes.

keepalive

Syntax **keepalive** *time-interval*
no keepalive

Context config>port>sonet-sdh>path>cisco-hdlc
config>port>tdm>ds1>channel-group>cisco-hdlc
config>port>tdm>ds3>cisco-hdlc
config>port>tdm>e1>channel-group>cisco-hdlc
config>port>tdm>e3>cisco-hdlc

Description This command specifies the interval, in seconds, used to send periodic keepalive packets. The receiver process expects to receive a keepalive packet every “keepalive interval”. The link is declared down if the receiver process does not receive a keepalive within the “timeout interval”. The link is declared up once the number of continual keepalive packets received equals to the up-count. The nodes at the two endpoints of the cHDLC link should be provisioned with the same values. A keepalive value of 0 means no keepalive packets are sent.

Default keepalive 10

Parameters *time-interval* — Specifies the interval used to send periodic keepalive packets.
Values 0 to 300

up-count

Syntax **up-count** *up-count*
no up-count

Context config>port>sonet-sdh>path>cisco-hdlc
config>port>tdm>ds1>channel-group>cisco-hdlc
config>port>tdm>ds3>cisco-hdlc
config>port>tdm>e1>channel-group>cisco-hdlc
config>port>tdm>e3>cisco-hdlc

Description This command configures the number of continual keepalive packets that have to be received in order to declare the link up. It is expected that the nodes at the two endpoints of the cHDLC link are provisioned with the same values.

Default up-count 1

Parameters	<i>up-count</i> — Specifies the number of continual keepalive packets that must be received in order to declare the link up.
Values	1 to 3

2.20.2.23 ATM Interface Commands

ATM Interface commands are supported on the 7750 SR only.

atm

Syntax	atm
Context	config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3 config>port>ml-bundle>ima
Description	This command enables the context to configure ATM interface properties.

cell-format

Syntax	cell-format {uni nni}
Context	config>port>sonet-sdh>path>atm config>port>tdm>ds1>channel-group>atm config>port>tdm>ds3>atm config>port>tdm>e1>channel-group>atm config>port>tdm>e3>atm config>port>ml-bundle>ima>atm
Description	This command configures the ATM cell format.
Parameters	uni — Specifies the user-to-network interface (UNI) cell format. nni — Specifies the network-to-network interface (NNI) cell format.

mapping

Syntax	mapping <i>mapping</i>
Context	config>port>tdm>ds3>atm

Description	This command configures the ATM cell mapping for DS-3 channels. The mapping value specifies the cell mapping that is to be used on this ATM interface.
Default	mapping direct
Parameters	<i>mapping</i> — The mapping value specifies the cell mapping that is to be used on this ATM interface.
Values	direct — Specifies direct cell mapping. plcp — Specifies PLCP cell mapping.

min-vp-vpi

Syntax	min-vp-vpi <i>value</i>
Context	config>port>sonet-sdh>path>atm config>port>ml-bundle>ima>atm config>port>tdm>ds1>channel-group>atm config>port>tdm>ds3>atm config>port>tdm>e1>channel-group>atm config>port>tdm>e3>atm
Description	This command sets the minimum allowable virtual path identifier (VPI) value that can be used on the ATM interface for a VPC.
Parameters	<i>value</i> — Specifies the minimum allowable VPI value that can be used on the ATM interface for a VPC.
Values	0 to 4095 (NNI) 0 to 255 (UNI)
Default	0

custom-buffer-mode

Syntax	[no] custom-buffer-mode
Context	config>port>sonet-sdh>path>atm
Description	This command configures ATM port custom buffer parameters.

buffer-pool

Syntax	buffer-pool <i>value</i>
Context	config>port>sonet-sdh>path>atm>custom-buffer-mode
Description	This command configures the ATM port buffer pool percentage.

Parameters	<i>value</i> — Specifies the percentage of the buffers configured on this interface over the maximum allowed.
Values	1 to 100

vc-threshold

Syntax	vc-threshold <i>buffer-threshold</i>
Context	config>port>sonet-sdh>path>atm>custom-buffer-mode
Description	This command configures the ATM port VC threshold.
Parameters	<i>buffer-threshold</i> — specifies the VC buffer threshold for all VCs configured on this port.
Values	190 to 117000

ilmi

Syntax	ilmi [<i>vpi/vci</i>] no ilmi						
Context	config>port>sonet-sdh>path>atm						
Description	This command creates an ILMI link PVCC by default on VPI/VC1 0/16. Deleting an ILMI link deletes the PVCC. ILMI is supported only on ATM interfaces on SONET/SDH paths.						
Parameters	<i>vpi/vci</i> — Specifies the PVC identifier (vpi/vci).						
Values	<table> <tr> <td>vpi</td><td>0 to 4095 (NNI)</td></tr> <tr> <td></td><td>0 to 255 (UNI)</td></tr> <tr> <td>vci</td><td>1, 2, 5 to 65535</td></tr> </table>	vpi	0 to 4095 (NNI)		0 to 255 (UNI)	vci	1, 2, 5 to 65535
vpi	0 to 4095 (NNI)						
	0 to 255 (UNI)						
vci	1, 2, 5 to 65535						

egress

Syntax	egress
Context	config>port>sonet-sdh>path>atm>ilmi
Description	This command enables the context to configure egress traffic attributes for the ILMI link.

ingress

Syntax	ingress

Context	config>port>sonet-sdh>path>atm>ilmi
Description	This command enables the context to configure ingress traffic attributes for the ILMI link.

traffic-desc

Syntax	traffic-desc <i>traffic-desc-profile-id</i> no traffic-desc
Context	config>port>sonet-sdh>path>atm>ilmi>egress config>port>sonet-sdh>path>atm>ilmi>ingress
Description	This command associates an ATM traffic descriptor profile to an ILMI link. It is recommended to change this to the traffic profile as defined in the ILMI specification.
Default	traffic-desc 1
Parameters	<i>traffic-desc-profile-id</i> — Specifies an existing ATM traffic descriptor profile. Traffic descriptor profiles are configured in the config>qos>atm-td-profile context. Values 1 to 1000

keep-alive

Syntax	keep-alive [poll-frequency <i>seconds</i>] [poll-count <i>value</i>] [test-frequency <i>seconds</i>] no keep-alive
Context	config>port>sonet-sdh>path>atm>ilmi
Description	This command configures keepalive parameters to monitor ILMI link connectivity. The no form of this command resets the default values on an ILMI link. Last Config Change: 03/29/2007 20:35:19 Poll Count:4 Poll Freq: 5 Test Freq: 1
Parameters	poll-frequency <i>seconds</i> — Specifies the amount of time, in seconds, between successive transmissions of ILMI messages on this interface for the purpose of detecting the establishment of ILMI connectivity. Values 1 to 255 <i>value</i> — Specifies the number of consecutive polls on this interface for which no ILMI response message is received before ILMI connectivity is declared lost. Values 1 to 255

test-frequency *seconds* — Specifies the frequency for testing for connectivity when the link is establishing before polling begins.

Values 0 to 255

protocol

Syntax	protocol <i>protocol-type</i> no protocol
Context	config>port>sonet-sdh>path>atm>ilmi
Description	This command configures the protocol.
Parameters	<i>protocol-type</i> — The <i>protocol-type</i> is an enumerated integer whose value indicates the ILM version of either 3.1 or 4.0 that is advertised by IME and also indicates the ILM type of either user-side or network-side. Values 4_0-user, 4_0-network. 3_1-user, 3_1-network

2.20.2.24 Frame Relay Commands

Frame Relay commands are supported on the 7750 SR only.

frame-relay

Syntax	frame-relay
Context	config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command allows access to the context to configure the Frame Relay Local Management Interface (LMI) operational parameters for a SONET/SDH PoS link, a DS-0 channel group, or a DS-3/E-3 port or channel. The port's mode must be set to access in config>port>sonet-sdh>path>mode access context. The port's encapsulation type must be set to frame-relay in the config>port>sonet-sdh>path>encap-type frame-relay context. The no form of this command removes the Frame Relay LMI operational parameters.

frf-12

Syntax	[no] frf-12
Context	config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay config>port>tdm>e3>frame-relay config>port>sonet-sdh>path>frame-relay
Description	This command defines the context to configure the parameters of FRF.12 Frame Relay fragmentation.

egress

Syntax	egress
Context	config>port>ml-bundle>mlfr config>port>sonet-sdh>path>frame-relay>frf-12 config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay>frf-12 config>port>tdm>e1>channel-group>frame-relay>frf-12 config>port>tdm>e3>frame-relay>frf-12
Description	This command enables the context to configure the egress QoS profile for an MLFR bundle or a Frame Relay port with FRF.12 UNI/NNI fragmentation enabled.

qos-profile

Syntax	qos-profile <i>profile-id</i> no qos-profile
Context	config>port>sonet-sdh>path>frame-relay>egress config>port>tdm>ds1>channel-group>frame-relay>frf-12>egress config>port>tdm>ds3>frame-relay>frf-12>egress config>port>tdm>e1>channel-group>frame-relay>frf-12>egress config>port>tdm>e3>frame-relay>frf-12>egress
Description	This command specifies the ingress or egress QoS profile to be used for the configuration of the egress QoS parameters of a Frame Relay port with FRF.12 UNI/NNI fragmentation enabled. The no form of the command removes the parameters from the configuration.
Parameters	<i>profile-id</i> — Specifies the profile number. The value can only be modified if the FR port is shut down. Values 1 to 65535

fragment-threshold

Syntax	fragment-threshold <i>threshold</i> no fragment-threshold
Context	config>port>sonet-sdh>path>frame-relay>frf-12 config>port>tdm>ds1>channel-group>frame-relay>frf-12>egress config>port>tdm>ds3>frame-relay>frf-12>egress config>port>tdm>e1>channel-group>frame-relay>frf-12>egress config>port>tdm>e3>frame-relay>frf-12>egress
Description	This command sets the maximum length in bytes of a fragment transmitted across a frame relay port with the FRF.12 UNI/NNI fragmentation enabled. The no form of this command resets the fragment threshold back to the default value.
Default	fragment-threshold 128
Parameters	<i>threshold</i> — Specifies the maximum fragment length, in bytes, to be transmitted across the FRF.12 port. Values 128 to 512

identifier

Syntax	identifier <i>frf16-link-id-string</i> no identifier
Context	config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>e1>channel-group>frame-relay
Description	This command defines the identifier for the FR bundle when used in an MLFR bundle. The no form of this command resets the value to null.
Default	null
Parameters	<i>frf16-link--id-string</i> — Specifies the bundle ID string up to 50 characters.

lmi-type

Syntax	lmi-type {ansi itu none rev1} no lmi-type
Context	config>port>multi-link-bundle>mlfr>frame-relay config>port>sonet-sdh>path>frame-relay config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay

config>port>tdm>e3>frame-relay

Description This command configures the Local Management Interface (LMI) type for Frame Relay interfaces. LMIs are sets of enhancements to the basic Frame Relay specification.

The **no** form of this command changes the LMI type back to the default value.

Default itu

Parameters **ansi** — Specifies to use ANSI T1.617 Annex D.
itu — Specifies to use ITU-T Q933 Annex A.
none — Specifies to disable Frame Relay LMI on the given port/channel.
rev1 — Specifies to use the Rev 1 version of ANSI T1.617 Annex D.

mode

Syntax **mode** {**dce** | **dte** | **bidir**}

Context config>port>multi-link-bundle>mlfr>frame-relay
config>port>sonet-sdh>path>frame-relay
config>port>tdm>ds1>channel-group>frame-relay
config>port>tdm>ds3>frame-relay
config>port>tdm>e1>channel-group>frame-relay
config>port>tdm>e3>frame-relay

Description This command sets the Frame Relay interface into the DCE, DTE, or Bidirectional mode of LMI operation. The DTE mode causes the router to send status inquiries over the interface. The DCE mode causes the router to respond to status inquiries. In bidirectional mode, the router performs both DTE and DCE operation over the FR interface. The bidirectional mode applies to the ANSI and ITU LMI types only.

This feature is used when two routers are connected back-to-back, running frame relay encapsulation.

Default mode dte

Parameters **dce** — Enables the DCE mode.
dte — Enables the DTE mode.
bidir — Enables the bidirectional mode for LMI types ANSI and ITU.

n391dte

Syntax **n391dte** *intervals*
no n391dte

Context config>port>multi-link-bundle>mlfr>frame-relay

```

config>port>sonet-sdh>path>frame-relay
config>port>tdm>ds1>channel-group>frame-relay
config>port>tdm>ds3>frame-relay
config>port>tdm>e1>channel-group>frame-relay
config>port>tdm>e3>frame-relay

```

Description This command sets the DTE full status polling interval for the Frame Relay Local Management Interface (LMI). The number specifies the frequency at which inquiries expect a full status report.

The **no** form of this command returns the **n391dte** counter to the default value.

Default n391dte 6

Parameters *intervals* — The number of exchanges to be done before requesting a full-status report. A value of 1 specifies to receive full-status messages only.

Values 1 to 255

n392dce

Syntax **n392dce** *threshold*
no n392dce

Context config>port>multi-link-bundle>mlfr>frame-relay
config>port>sonet-sdh>path>frame-relay
config>port>tdm>ds1>channel-group>frame-relay
config>port>tdm>ds3>frame-relay
config>port>tdm>e1>channel-group>frame-relay
config>port>tdm>e3>frame-relay

Description This command sets the DCE error threshold for the Frame Relay Local Management Interface (LMI).

The threshold specifies the number of errors needed to bring down a link.

The **no** form of this command returns the **n392dce** counter to the default value.

Default n392dce 3

Parameters *threshold* — Specifies the number of errors that will place the channel in an operationally down state.

Values 1 to 10

n392dte

Syntax **n392dte** *threshold*
no n392dte

Context	<pre>config>port>multi-link-bundle>mlfr>frame-relay config>port>sonet-sdh>path>frame-relay config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay config>port>tdm>e3>frame-relay</pre>
Description	This command sets the DTE error threshold for the Frame Relay Local Management Interface (LMI). The count specifies the number of errors needed to bring down a link. The no form of this command returns the n392dte counter to the default value.
Default	n392dte 3
Parameters	<i>threshold</i> — Specifies the number of errors that will place the path or channel in an operationally down state. Values 1 to 10

n393dce

Syntax	<pre>n393dce <i>count</i> no n393dce</pre>
Context	<pre>config>port>multi-link-bundle>mlfr>frame-relay config>port>sonet-sdh>path>frame-relay config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay config>port>tdm>e3>frame-relay</pre>
Description	This command sets the DCE monitored event count for the Frame Relay Local Management Interface (LMI). The no form of this command returns the n393dce counter to the default value.
Default	n393dce 4
Parameters	<i>count</i> — Specifies the diagnostic window used to verify link integrity on the DCE interface. Values 1 to 10

n393dte

Syntax	<pre>n393dte <i>count</i> no n393dte</pre>

Context	<pre>config>port>sonet-sdh>path>frame-relay config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay config>port>tdm>e3>frame-relay</pre>
Description	This command sets the DTE monitored event count for the Frame Relay Local Management Interface (LMI). The no form of this command returns the n393dte counter to the default value.
Default	n393dte 4
Parameters	<i>count</i> — Specifies the diagnostic window used to verify link integrity on the DTE interface. Values 1 to 10

t391dte

Syntax	<pre>t391dte keepalive no t391dte</pre>
Context	<pre>config>port>sonet-sdh>path>frame-relay config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay config>port>tdm>e3>frame-relay</pre>
Description	This command sets the DTE keepalive timer for the Frame Relay Local Management Interface (LMI). This number specifies the period at which the DTE sends out a keepalive response request to the DCE and updates status depending on the DTE error threshold value. The no form of this command returns the t391dte keepalive timer to the default value.
Default	t391dte 10
Parameters	<i>keepalive</i> — Specifies the interval in seconds between status inquiries issued by the DTE. Values 5 to 30

t392dce

Syntax	<pre>t392dce keepalive no t392dce</pre>

Context	<pre>config>port>sonet-sdh>path>frame-relay config>port>tdm>ds1>channel-group>frame-relay config>port>tdm>ds3>frame-relay config>port>tdm>e1>channel-group>frame-relay config>port>tdm>e3>frame-relay</pre>
Description	This command sets the DCE keepalive timer for the Frame Relay Local Management Interface (LMI). This number specifies the period at which the DCE checks for keepalive responses from the DTE and updates status depending on the DCE error threshold value. The no form of this command returns the t392dce keepalive timer to the default value.
Default	t392dce 15
Parameters	<i>keepalive</i> — Specifies the expected interval in seconds between status inquiries issued by the DTE equipment. Values 5 to 30

2.20.2.25 TDM Commands

TDM commands are only supported on the 7750 SR.

tdm

Syntax	tdm
Context	config>port
Description	This command enables the context to configure DS-1/E-1 and DS-3/E-3 parameters for a port on a channelized MDA T1/E1. This context cannot be accessed on non-channelized MDAs. TDM is a mechanism to divide the bandwidth of a stream into separate channels or time slots by assigning each stream a different time slot in a set. TDM repeatedly transmits a fixed sequence of time slots over a single transmission channel. Each individual data stream is reassembled at the receiving end based on the timing.

ds1

Syntax	[no] ds1 <i>ds1-id</i>
Context	config>port>tdm

Description	This command enables the context to configure digital signal level 1 (DS-1) frame parameters. The T-Carrier system was the first successful system that supported digitized voice transmission. The original transmission rate (1.544 Mb/s) in the T-1 (DS-1) line is commonly used by Internet service providers (ISPs) to connect to the Internet. North America uses the T-Carrier system while Europe uses the E-Carrier system of transmission, using multiples of the DS- system. Digital signals are carried inside the carrier systems. T-1 transmits DS-1-formatted data at 1.544 Mb/s through the network. The corresponding European carrier is E-1 with a data rate of 2.048 Mb/s. E-1 and T-1 (DS-1) can be interconnected for international use. The no form of this command disables DS-1 capabilities.
Parameters	<i>ds1-id</i> — Identifies the DS-1 channel being created. Values DS1: 1 to 28, ds1-sonet-sdh-index

ds3

Syntax	[no] ds3 [<i>sonet-sdh-index</i>]
Context	config>port>tdm
Description	This command enables the context to configure DS-3 parameters. DS-3 lines provide a speed of 44.736 Mb/s and is also frequently used by service providers. DS-3 lines carry 28 DS-1 signals and a 44.736 Mb/s data rate. A DS-3 connection typically supports data rates of about 43 Mb/s. A T-3 line actually consists of 672 individual channels, each supporting 64 kb/s. T-3 lines are used mainly by Service Providers to connect to the Internet backbone and for the backbone itself. Depending on the MDA type, the DS-3 parameters must be disabled if clear channel is enabled by default (for example, on the m12-ds3 MDA). Clear channel is a channel that uses out-of-band signaling, not in-band signaling, so the channel's entire bit rate is available. Channelization must be explicitly specified. Note that if DS-3 nodes are provisioned on a channelized SONET/SDH MDA you must provision the parent STS-1 SONET/STM0 SDH path first. North America uses the T-Carrier system while Europe uses the E-Carrier system of transmission, using multiples of the DS system. Digital signals are carried inside the carrier systems. The no form of this command disables DS-3 capabilities.
Parameters	<i>sonet-sdh-index</i> — Specifies the components making up the specified SONET/SDH Path. Depending on the type of SONET/SDH port the <i>sonet-sdh-index</i> must specify more path indexes to specify the payload location of the path. The <i>sonet-sdh-index</i> differs for SONET and SDH ports.

e1

Syntax	e1 [<i>e1-id</i>]
Context	config>port>tdm
Description	This command enables the context to configure E-1 parameters. E-1 is a basic time division multiplexing scheme used to carry digital circuits. It is also a standard WAN digital communication format designed to operate over copper facilities at a rate of 2.048 Mb/s. North America uses the T-Carrier system while Europe uses the E-Carrier system of transmission, using multiples of the DS system. Digital signals are carried inside the carrier systems. The no form of this command disables E-1 capabilities.
Parameters	<i>e1-id</i> — Specifies the E-1 channel being created.
Values	E1: 1 to 21, e1-sonet-sdh-index

e3

Syntax	[no] e3 [<i>sonet-sdh-index</i>]
Context	config>port>tdm
Description	This command enables the context to configure E-3 parameters. E-3 lines provide a speed of 44.736 Mb/s and is also frequently used by service providers. E-3 lines carry 16 E-1 signals with a data rate of 34.368 Mb/s. An E-3 connection typically supports data rates of about 43 Mb/s. An E-3 line actually consists of 672 individual channels, each supporting 64 kb/s. E-3 lines are used mainly by Service Providers to connect to the Internet backbone and for the backbone itself. Depending on the MDA type, the E-3 parameters must be disabled if clear channel is enabled by default (for example, on the m12-ds3e3 MDA). Clear channel is a channel that uses out-of-band signaling, not in-band signaling, so the channel's entire bit rate is available. Channelization must be explicitly specified. Note that if E-3 nodes are provisioned on the channelized SONET/SDH MDA you must provision the parent STS-1 SONET/STM0 SDH path first. North America uses the T-Carrier system while Europe uses the E-Carrier system of transmission, using multiples of the DS system. Digital signals are carried inside the carrier systems. The no form of this command disables E-3 capabilities.

Parameters *sonet-sdh-index* — Specifies the components making up the specified SONET/SDH Path. Depending on the type of SONET/SDH port the *sonet-sdh-index* must specify more path indexes to specify the payload location of the path. The *sonet-sdh-index* differs for SONET and SDH ports.

bert

Syntax **bert** {**2e3** | **2e9** | **2e11** | **2e15** | **2e20** | **2e20q** | **2e23** | **ones** | **zeros** | **alternating**} *duration*
no bert

Context config>port>tdm>ds1
config>port>tdm>ds3
config>port>tdm>e1
config>port>tdm>e3

Description This command initiates or restarts a Bit Error Rate Test (BERT) on the associated DS-1/E-1 or DS-3/E-3 circuit.

The associated DS-1, E-1, DS-3, or E-3 must be in a shutdown (admin down) state to initiate this test.

The **no** form of the command terminates the BERT test if it has not yet completed.

Notes:

- This command is not saved in the router configuration between boots.
- The 4-port OC-3/STM-1 and the 1-port OC-12/STM-4 ASAP MDA supports up to 28 concurrent BERT tests per MDA. The 4-port and 12-port DS-3/E-3 ASAP MDAs support a single BERT test per MDA. An attempt to configure more BERT tests can result in an error indicating an operation failure due to resource exhaustion.
- If the ASAP MDA BERT error insertion **rate** command is executed when tests are running, it will not take effect until test is restarted.

Default bert 2e3

Parameters *duration* — Sets the duration for the BERT test.

Values Up to 24 hours, in seconds or hh:mm:ss format

ones — Sends an all ones pattern.

zeros — Sends an all zeros pattern.

alternating — Sends an alternating ones and zeros pattern.

2e3 — Sends a pseudo-random $2^3 - 1$ pattern.

2e9 — Sends a pseudo-random $2^9 - 1$ pattern.

2e15 — Sends a pseudo-random $2^{15} - 1$ pattern.

2e20 — Sends a pseudo-random $2^{20} - 1$ pattern. Not available on channelized ASAP MDAs.

2e23 — Sends a pseudo-random $2^{23} - 1$ pattern.

bit-error-insertion

Syntax	bit-error-insertion <i>rate</i> no bit-error-insertion
Context	config>port>tdm>ds1 config>port>tdm>ds3 config>port>tdm>e1 config>port>tdm>e3
Description	This command inserts bit errors into a running BERT test. The number of errors inserted corresponds to $10^{(-rate)}$. A rate of 0 will cause 1 error in every bit transmitted. A rate of 7 will cause an error rate of $10^{(-7)}$, or 1 error in every one billion bits transmitted. The no command disables the insertion of bit errors into the bit error rate test stream. Note that this command is not saved in the router configuration between boots.
Default	no bit-error-insertion
Parameters	<i>rate</i> — Specifies the bit error rate, expressed as an integer. Values 2 to 7

buildout

Syntax	buildout { long short }
Context	config>port>tdm
Description	This command specifies line buildout (cable length) for physical DS-1/DS-3 ports.
Default	buildout short
Parameters	long — Sets the line buildout for length runs up to 450 feet. short — Sets the line buildout for length runs up to 225 feet.

hold-time

Syntax	hold-time {[up <i>hold-time-up</i>] [down <i>hold-time-down</i>]} no hold-time

Context	config>port>tdm
Description	This command configures link dampening timers in 100s of milliseconds. This guards against reporting excessive interface transitions. This is implemented by not advertising subsequent transitions of the interface to upper layer protocols until the configured timer has expired. This command is only supported on the m4-chds3-as, m12-chds3-as, and c4-ds3 MDAs.
Default	no hold-time
Parameters	<i>hold-time-up</i> — Configures the hold-timer for link up event dampening. A value of zero (0) indicates that an up transition is reported immediately. Values 0 to 100 in 100s of milliseconds (default 0) <i>hold-time-down</i> — The hold-timer for link down event dampening. A value of zero (0) indicates that a down transition is reported immediately. Values 0 to 100 in 100s of milliseconds (default 5)

lbo

Syntax	lbo [0dB -7.5dB -15.0dB -22.5dB]								
Context	config>port>tdm								
Description	This command applies only to a DS-1 port configured with a 'long' buildout (see the buildout command). Specify the number of decibels the transmission signal decreases over the line. For 'short' buildout the following values are valid: lboNotApplicable — Not applicable For 'long' buildout the following values are valid: <table><tr><td>lbo0dB</td><td>For 0 dB</td></tr><tr><td>lboNeg7p5dB</td><td>For -7.5 dB</td></tr><tr><td>lboNeg15p0dB</td><td>For -15.0 dB</td></tr><tr><td>lboNeg22p5dB</td><td>For -22.5 dB</td></tr></table>	lbo0dB	For 0 dB	lboNeg7p5dB	For -7.5 dB	lboNeg15p0dB	For -15.0 dB	lboNeg22p5dB	For -22.5 dB
lbo0dB	For 0 dB								
lboNeg7p5dB	For -7.5 dB								
lboNeg15p0dB	For -15.0 dB								
lboNeg22p5dB	For -22.5 dB								

The default for 'short' build out is 'NotApplicable' while the default for 'long' buildout is 'lbo0dB'.

length

Syntax	length {133 266 399 533 655}
Context	config>port>tdm

Description This command applies only to a DS-1 port configured with a 'short' buildout. The **length** command configures the length of the line (in feet). For line lengths longer than 655 feet, configure the DS-1 port buildout as 'long'.

For 'long' buildout the following values are valid:

NotApplicable — Not applicable

For 'short' buildout the following values are valid:

- 0 to 133 For line length from 0 to 133 feet
- 134 to 266 For line length from 134 to 266 feet
- 267 to 399 For line length from 267 to 399 feet
- 400 to 533 For line length from 400 to 533 feet
- 534 to 655 For line length from 534 to 655 feet

The default for 'long' buildout is 'NotApplicable' while the default for 'short' buildout is '0 to 133'.

channel-group

Syntax **[no]** **channel-group** *channel-group-id*

Context config>port>tdm>ds1>channel-group
config>port>tdm>ds3>channel-group
config>port>tdm>e1>channel-group

Description This command creates DS0 channel groups in a channelized DS1 or E1 circuit. Channel groups cannot be further subdivided.

The **no** form of this command deletes the specified DS1 or E1 channel.

Parameters *channel-group-id* — Identifies the channel-group ID number.

Values DS1: 1 to 24
E1: 1 to 32

channelized

Syntax **channelized** {**ds1** | **e1**}
no channelized

Context config>port>tdm>ds3

Description	This command specifies that the associated DS-3 is a channelized DS-3 with DS-1/E-1 sub-channels. Depending on the MDA type, the DS-3 parameters must be disabled if clear channel is the default (for example, on m12-ds3 MDAs). Clear channel is a channel that uses out-of-band signaling, not in-band signaling, so the channel's entire bit rate is available. Channelization must be explicitly specified. The no form specifies the associated DS-3 is a clear channel circuit and cannot contain sub-channel DS-1s/E-1s. The sub-channels must be deleted first before the no command is executed.
Default	no channelized.
Parameters	ds1 — Specifies that the channel is DS-1. e1 — Specifies that the channel is E-1.

cisco-hdlc

Syntax	cisco-hdlc
Context	<pre>config>port>sonet-sdh>path config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3</pre>
Description	This command enables the context to configure Cisco HDLC parameters. Cisco HDLC is an encapsulation protocol that governs information transfer. It specifies a data encapsulation method on synchronous serial links using frame characters and checksums. Cisco HDLC monitors line status on a serial interface by exchanging keepalive request messages with peer network devices. It also allows routers to discover IP addresses of neighbors by exchanging Serial Link Address Resolution Protocol (SLARP) address-request and address-response messages with peer network. Only IES SAPs (including SAPs in VPRN service) can provision a Cisco-HDLC-capable configuration.

clock-source

Syntax	clock-source {loop-timed node-timed adaptive differential}
Context	<pre>config>port>tdm>ds1 config>port>tdm>ds3 config>port>tdm>e1 config>port>tdm>e3</pre>

Description This command configures the clock to be used for transmission of data out towards the line. The options are to use the locally recovered clock from the line's receive data stream, the node central reference, or an adaptively recovered clock using the received packets.

The following tables show MDAs that support loop timing at DS3/E3 and DS1/E1 channelization options.

TDM DS3/E3	LoopTimed	Default
Channelized OC-12	No	node-timed
Channelized OC-3	No	node-timed
Channelized DS3/E3	No	node-timed
Channelized ASAP OC-12	Yes	node-timed
Channelized ASAP OC-3	Yes	node-timed
Channelized ASAP DS3/E3	Yes	node-timed
CES OC-3	Yes	node-timed

TDM DS1/E1	LoopTimed	Default
Channelized OC-12	Yes	loop-timed
Channelized OC-3	Yes	loop-timed
Channelized DS3/E3	Yes	loop-timed
Channelized ASAP OC-12	Yes	loop-timed
Channelized ASAP OC-3	Yes	loop-timed
Channelized ASAP DS3/E3	Yes	loop-timed
CES OC-3	Yes	loop-timed

Parameters

loop-timed — The link recovers the clock from the received data stream.

node-timed — The link uses the internal clock when transmitting data.

adaptive — The clock is adaptively recovered from the rate at which data is received and not from the physical layer. Adaptive timing is only supported on ds1 and e1 channels.

differential — The clock is recovered from differential RTP timestamp header. Differential timing is only supported on ds1 and e1 channels.

crc

Syntax **crc {16 | 32}**

Context

```
config>port>tdm>ds1>channel-group
config>port>tdm>ds3
config>port>tdm>e1>channel-group
config>port>tdm>e3
```

Description	This command configures the precision of the cyclic redundancy check (CRC).
Default	crc 16 for non-ATM channel groups configured under DS-1, E-1 and for non-ATM E-3 and DS-3 channel/ports. crc 32 for ATM channel-groups configured under DS-1 and E-1, and for ATM E-3 and DS-3 channels/ports. The default cannot be changed.
Parameters	16 — Uses 16 bit checksum for the associated port/channel. 32 — Uses 32 bit checksum for the associated port/channel.

down-count

Syntax	down-count <i>down-count</i> no down-count
Context	config>port>sonet-sdh>path>cisco-hdlc config>port>tdm>ds1>channel-group>cisco-hdlc config>port>tdm>ds3>cisco-hdlc config>port>tdm>e1>channel-group>cisco-hdlc config>port>tdm>e3>cisco-hdlc
Description	This command configures the number of keepalive intervals that must pass without receiving a keepalive packet before the link is declared down. It is expected that the nodes at the two endpoints of the cHDLc link are provisioned with the same values.
Default	down-count 3
Parameters	<i>down-count</i> — Specifies the number of keep alive intervals that must pass without receiving a keep alive packet before the link is declared down. Values 3 to 16

encap-type

Syntax	encap-type {atm bcp-null bcp-dot1q ipcp ppp-auto frame-relay wan-mirror cisco-hdlc cem}
Context	config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command configures the encapsulation method used to on the specified port, path, or channel. This parameter can be set on both access and network ports.

When the **encap-type** is set to ATM the CRC, timeslots, scrambling (if applicable), and idle-cycle-flags are set to ATM defaults respectively. When the encap-type is changed from ATM, those parameters are set to their non-ATM defaults.

When the **encap-type** is ATM, ATM sub-layer verification (GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs)*) is automatically enabled. When ATM PLCP cell mapping is used, the results of this verification include:

- PLCP Severely Errored Framing Seconds
- PLCP Alarm State
- PLCP Unavailable Seconds Counter

When ATM direct cell mapping is used, the result of the verification includes:

- Out of Cell Delineation (OCD) event count. The OCD event count is described in RFC 2515, *Definitions of Managed Objects for ATM Management*. Note that multiple events occurring within a second will be counted as 1 event for ASAP MDAs as a result of a hardware limit.
- Loss of Cell Delineation defect/alarm. The LCD defect/alarm is defined in RFC 2515, *Definitions of Managed Objects for ATM Management*. When a path is in an LCD defect state, the path's operational status will be down. When a path exits the LCD state, the path's operational status will change to up (assuming nothing else causes the path to stay down). A trap is raised to indicate the LCD status change. Also, a P-RDI is sent to indicate the defect to the remote end.

The **no** form of this command restores the default.

Default	encap-type bcp-null
Parameters	atm — Specifies the encapsulation on the port is ATM. bcp-null — When selected, this keyword specifies that only a single service is configured on this channel and IEEE 802.1Q tags are not used as a service delimiter. Any IEEE 802.1Q tags encountered are regarded as part of the customer payload and transparently forwarded. When bcp-null encapsulation is specified, the PPP Bridge Control Protocol (BCP) is activated and all packets on this access port will be encapsulated in accordance with the BCP protocol. bcp-dot1q — When selected, this keyword specifies that ingress frames carry IEEE 802.1Q tags and the tags are used as service delimiter. Any untagged packets are silently discarded with exception of protocol specific packets. When bcp-dot1q encapsulation is specified, the PPP Bridge Control Protocol (BCP) is activated and all packets on this access port will be encapsulated in accordance with the BCP protocol. ipcp — Ingress frames are encapsulated according to the IP Control Protocol. When ipcp encapsulation is specified, the PPP IP Control Protocol will be activated and only packets that comply with IPCP encapsulation are processed; others are silently discarded.

ppp-auto — (Network mode) Enables PPP on the associated port/channel. The activation of ipcp and mplscp is automatically enabled depending on the protocol configuration. This encaps type is only valid on ports/channels in network mode.

frame-relay — Enables frame relay on the associated port/channel.

wan-mirror — The port is used for mirroring of frame-relay and POS ports. On these ports, no link management protocol will run.

cisco-hdlc — Monitors line status on a serial interface by exchanging keepalive request messages with peer network devices.

cem — Specifies that on circuit emulation CMAs and MDAs, only the **cem** encaps-type is supported. All other values are blocked with an appropriate warning. The **cem** encaps-type is not supported on other CMAs and MDAs and are blocked with an appropriate warning.

feac-loop-respond

Syntax	[no] feac-loop-respond
Context	config>port>tdm>ds3 config>port>tdm>e3
Description	This command enables the associated DS-3 interface to respond to remote loop signals. The DS-3 far-end alarm and control (FEAC) signal is used to send alarm or status information from the far-end terminal back to the local terminal. DS-3 loopbacks at the far-end terminal from the local terminal are initiated. The no form of this command prevents the associated DS-3/E-3 interface from responding to remote loop signals.
Default	no feac-loop-respond

framing (DS-1)

Syntax	framing {esf sf unframed-ds1}
Context	config>port>tdm>ds1
Description	This command specifies the DS-1 framing to be used with the associated channel.
Default	framing esf
Parameters	esf — Configures the DS-1 port for extended super frame framing. sf — Configures the DS-1 port for super frame framing.

unframed-ds1 — Specifies ds-1 unframed (G.703) mode for DS-1 interfaces. This parameter allows the configuration of an unstructured DS-1 channel on a CES MDA. In G.704, timeslot 0 is used to carry timing information by a service provider, thus, only 31 slots are made available to the end user. In G.703, all 32 time slots are available to the end user. Timing is provided by the end user. When an e1-unframed channel is shutdown, it sends the AIS pattern to the far-end DS-1 which does not react. The operational status remains up and no alarms are generated while the near-end (shutdown) is operationally down. This is normal behavior since the G.703 option does not have framing. G.703 framing is only applicable for FR, PPP, and cHDLC encapsulations.

framing (E-1)

Syntax	framing {no-crc-g704 g704 e1-unframed}
Context	config>port>tdm>e1
Description	This command specifies the E-1 framing to be used with the associated channel.
Default	framing g704
Parameters	g704 — Configures the E-1 port for G.704 framing. no-crc-g70 — Configures the E-1 for G.704 with no CRC4. e1-unframed — Specifies E-1 unframed (G.703) mode for E-1 interfaces. This parameter also allows the configuration of an unstructured E-1 channel on an ASAP or CES MDA. In G.704, timeslot 0 is used to carry timing information by a service provider, thus, only 31 slots are made available to the end user. In G.703, all 32 time slots are available to the end user. Timing is provided by the end user. When an e1-unframed channel is shutdown, it sends the AIS pattern to the far-end E-1 which does not react. The operational status remains up and no alarms are generated while the near-end (shutdown) is operationally down. This is normal behavior since the G.703 option does not have framing. G.703 framing is only applicable for FR, PPP, and cHDLC and CEM encapsulations.

framing (DS3)

Syntax	framing {c-bit m23 unframed-ds3}
Context	config>port>tdm>ds3
Description	This command specifies DS-3 framing for the associated DS-3 port or channel.
Default	framing c-bit
Parameters	c-bit — Configures the DS-3 port/channels for C-Bit framing. m23 — Configures the DS-3 port/channel for M23 framing.

unframed-ds3 — Specifies ds-3 unframed mode for DS-3 interfaces. This parameter allows the configuration of an unstructured DS-3 channel on a CES MDA.

framing (E-3)

Syntax	framing {g751 g832 unframed-e3}
Context	config>port>tdm>e3
Description	This command specifies E-3 framing for the associated E-3 port or channel.
Default	for E-3 non-ATM: framing g751 and cannot be changed. for E-3 ATM: framing g832 and cannot be changed.
Parameters	g751 — Configures the E-3 port/channel for g751 framing. g832 — Configures the E-3 port/channel for g832 framing. unframed-e3 — Specifies e-3 unframed mode for E-3 interfaces. This parameter allows the configuration of an unstructured E-3 channel on a CES MDA.

idle-cycle-flag

Syntax	idle-cycle-flag {flags ones} no idle-cycle-flag
Context	config>port>tdm>ds1>channel-group config>port>tdm>ds3 config>port>tdm>e1>channel-group config>port>tdm>e3
Description	This command configures the value that the HDLC TDM DS-0, E-1, E-3, DS-1, or DS-3 interface transmits during idle cycles. For ATM ports/channels/channel-groups, the configuration does not apply and only the no form is accepted. The no form of this command reverts the idle cycle flag to the default value.
Default	flags (0x7E) no flags (ATM)
Parameters	flags — Specifies that 0x7E is used as the idle cycle flag. ones — Specifies that 0xFF is used as the idle cycle flag.

idle-payload-fill

Syntax	idle-payload-fill {all-ones}

idle-payload-fill *pattern pattern*
no idle-payload-fill

Context	config>port>tdm>ds1>channel-group config>port>tdm>e1>channel-group
Description	This command defines the data pattern to be transmitted when the circuit emulation service is not operational or temporarily experiences under-run conditions. This command is only valid for cesopsn and cesopsn-cas circuit emulation services. It is blocked with a warning for unstructured (satop) circuit emulation services.
Default	idle-payload-fill all-ones
Parameters	all-ones — Defines the 8 bit value to be transmitted as 11111111. <i>pattern</i> — Transmits a user-defined pattern. Values 0 to 255, accepted in decimal, hex or binary

idle-signal-fill

Syntax	idle-signal-fill {all-ones} idle-signal-fill pattern pattern no idle-signal-fill
Context	config>port>tdm>ds1>channel-group config>port>tdm>e1>channel-group
Description	This command defines the signaling pattern to be transmitted when the circuit emulation service is not operational or temporarily experiences under-run conditions. This command is only valid for cesopsn-cas circuit emulation services. It is blocked with a warning for unstructured (satop) and basic cesopsn circuit emulation services.
Default	idle-signal-fill all-ones
Parameters	all-ones — Defines the 8 bit value to be transmitted as 11111111. <i>pattern</i> — Transmits a user-defined pattern. Values 0 to 15, accepted in decimal, hex or binary

invert-data

Syntax	[no] invert-data
Context	config>port>tdm>ds1 config>port>tdm>e1
Description	This command causes all data bits to be inverted, to guarantee ones density. Typically used with AMI line encoding.

Default no invert-data

loopback

Syntax	loopback { line internal fdl-ansi fdl-bellcore payload-ansi inband-ansi inband-bellcore } no loopback
Context	config>port>tdm>ds1
Description	This command puts the specified port or channel into a loopback mode. The corresponding port or channel must be in a shutdown state in order for the loopback mode to be enabled. The upper level port or channel or parallel channels should not be affected by the loopback mode Note that this command is not saved in the router configuration between boots. The no form of this command disables the specified type of loopback.
Default	no loopback
Parameters	line — Places the associated port or channel into a line loopback mode. A line loopback loops frames received on the corresponding port or channels back to the remote router. internal — Places the associated port or channel into an internal loopback mode. An internal loopback loops the frames from the local router back at the framer. fdl-ansi — Requests FDL line loopback according to ANSI T1.403. fdl-bellcore — Requests FDL line loopback according to Bellcore TR-TSY-000312. payload-ansi — Requests payload loopback using ANSI signaling. inband-ansi — Requests inband line loopback according to ANSI T1.403. inband-bellcore — Requests inband line loopback according to Bellcore signaling.

loopback

Syntax	loopback { line internal remote } no loopback
Context	config>port>tdm>ds3
Description	This command puts the specified port or channel into a loopback mode. The corresponding port or channel must be in a shutdown state in order for the loopback mode to be enabled. The upper level port or channel or parallel channels should not be affected by the loopback mode.

Note that this command is not saved in the router configuration between boots.

The **no** form of this command disables the specified type of loopback.

Default	no loopback
Parameters	line — Places the associated port or channel into a line loopback mode. A line loopback loops frames received on the corresponding port or channels back to the remote router. internal — Places the associated port or channel into an internal loopback mode. A internal loopback loops the frames from the local router back at the framer. remote — Sends a signal to the remote device to provide a line loopback.

loopback

Syntax	loopback {line internal} no loopback
Context	config>port>tdm>e1 config>port>tdm>e3
Description	This command puts the specified port or channel into a loopback mode. The corresponding port or channel must be in a shutdown state in order for the loopback mode to be enabled. The upper level port or channel or parallel channels should not be affected by the loopback mode Note that this command is not saved in the router configuration between boots. The no form of this command disables the specified type of loopback.
Default	no loopback
Parameters	line — Places the associated port or channel into a line loopback mode. A line loopback loops frames received on the corresponding port or channels back to the remote router. internal — Places the associated port or channel into an internal loopback mode. An internal loopback loops the frames from the local router back at the framer.

mdl

Syntax	mdl {eic lic fic unit pfi port gen} mdl-string no mdl [eic lic fic unit pfi port gen]
Context	config>port>tdm>ds3
Description	This command configures the maintenance data link (MDL) message for a DS-3/E-3.

This command is only applicable if the DS-3/E-3 is using C-bit framing (see the [framing \(DS3\)](#) command).

The **no** form of this command removes the MDL string association and stops the transmission of any IDs.

Default no mdl

Parameters *mdl-string* — Specifies an MDL message up to 38 characters long on a DS-3.
eic — Specifies the equipment ID code up to 10 characters long.
lic — Specifies the equipment ID code up to 11 characters long.
fic — Specifies the ID code up to 10 characters long.
unit — Specifies the unit ID code up to 6 characters long.
pfi — Specifies the facility ID code up to 38 characters long.
port — Specifies the port ID code up to 38 characters long.
gen — Specifies the generator number to send in the MDL test signal message up to 38 characters long.

mdl-transmit

Syntax **mdl-transmit** {*path* | *idle-signal* | *test-signal*}
no mdl-transmit [*path* | *idle-signal* | *test-signal*]

Context config>port>tdm>ds3

Description This command enables the transmission of an MDL message on a DS-3/E-3 over channelized interface.

The **no** form of this command disables transmission of the specified message or all messages.

Default no mdl-transmit

Parameters **path** — Specifies the MDL path message.
idle-signal — Specifies the MDL idle signal message.
test-signal — Specifies the MDL test signal message.

national-bits

Syntax **national-bits** *sa4 sa5 sa6 sa7 sa8*
no national-bits

Context config>port>tdm>e1

Description This command configures the national use bits.

Parameters *sa-bits* — Disables or enables SA bits.

Values 0, 1

remote-loop-respond

Syntax **[no] remote-loop-respond**

Context config>port>tdm>ds1

Description When enabled, the channel responds to requests for remote loopbacks.

Default no remote-loop-respond — The port will not respond.

report-alarm

Syntax **[no] report-alarm [ais] [los] [oof] [rai] [looped] [ber-sd] [ber-sf]**

Context config>port>tdm>ds1
config>port>tdm>e1

Description This command enables logging of DS-1/DS-3 or E-1/E-3 alarms for DS-1/DS-3 or E-1/E-3 ports or channels.

The **no** form of this command disables logging of the specified alarms.

Parameters **ais** — Reports alarm indication signal errors. When configured, **ais** alarms are not raised and cleared.

Default **ais** alarms are issued

los — Reports loss of signal errors. When configured, **los** traps are not raised and cleared.

Default **los** traps are issued.

oof — Reports out-of-frame errors. When configured, **oof** alarms are not raised and cleared.

Default **oof** alarms are not issued.

rai — Reports resource availability indicator events. When configured, **rai** events are not raised and cleared.

Default **rai** alarms are not issued

looped — Reports looped packets errors.

looped alarms are not issued**lof** — Reports loss of frame errors. When configured, **lof** traps are not raised and cleared.

Default **lof** traps are issued

ber-sd — Specifies the BER that specifies signal degradation.

ber-sf — Specifies the BER that specifies signal failure.

signal-mode

Syntax	signal-mode {cas} no signal-mode
Context	config>port>tdm>ds1 config>port>tdm>e1
Description	This command activates the signal mode on the channel. When enabled, it uses routing information to direct the payload of voice or data to its destination. The no form of the command reverts to the default value.
Default	no signal-mode
Parameters	cas — Specifies channel associated signaling.

speed

Syntax	speed {56 64}
Context	config>port>tdm>ds1>channel-group config>port>tdm>e1>channel-group
Description	This command sets the speed of the DS-0 channels used in the associated channel-group.
Default	speed 64
Parameters	56 — Specifies that 56k byte (7-bits per byte) encoding will be used for the associated DS-0 channels. This channel speed value is only supported on the m4-chds3-as and m12-chds3-as MDAs and on DS-1 channels (ESF and SF framing) and not on E-1 (G.704) channels. 64 — Specifies that 64k byte (8-bits per byte) encoding will be used for the associated DS-0 channels.

subrate

Syntax	subrate {digital-link larscom} rate-step

no subrate

Context	config>port>tdm>ds3
Description	This command configures the channel service unit (CSU) compatibility mode to interoperate with existing DS-3 subrate standards. This configuration applies only for non-channelized DS-3s on ASAP TDM MDAs. The no form of this command remove the subrate functionality.
Default	no subrate
Parameters	digital-link — Enables the Digital-Link (Quick Eagle) CSU compatibility mode. larscom — Enables the Larscom CSU compatibility mode. rate-step — Specifies the subrate value for the associated DS-3.
Values	1 to 147 (digital-link) 1 to 14 (larscom)

threshold

Syntax	threshold {ber-sd ber-sf} rate {1 5 10 50 100} no threshold {ber-sd ber-sf}
Context	config>port>tdm>ds1 config>port>tdm>e1
Description	This command configures the line signal degradation bit error rate (BER) and line signal failure thresholds. Line signal (b2) bit interleaved parity error rates are measured and when they cross either the degradation or failure thresholds alarms are raised (see the report-alarm command), furthermore if the failure threshold is crossed the link will be set to operationally down. The no form of this command reverts to the default value.
Default	threshold ber-sd rate 5 threshold ber-sf rate 50
Parameters	ber-sd — Specifies the BER that specifies signal degradation. ber-sf — Specifies the BER that specifies signal failure. rate — Specifies the number of errors, in millions.

timeslots

Syntax	timeslots <i>timeslots</i>

	no timeslots
Context	config>port>tdm>ds1>channel-group config>port>tdm>e1>channel-group
Description	This command defines the list of DS-0 timeslots to be used in the DS-1 or E-1 channel-group. The timeslots are defaulted as defined below when encap-type is set to/from atm. ATM channel groups do not allow timeslots to change. The no form of this command removes DS-0 timeslots from a channel group.
Parameters	<i>timeslots</i> — Specifies the timeslot(s) to be associated with the channel group. The value can consist of a list of timeslots. Each member of the list can either be a single timeslot or a range of timeslots. Values 1 to 24 for DS-1 interfaces (the full range is auto-configured for ATM channel groups and cannot be changed) 2 to 32 for E-1 interfaces (the 2 to 16, 18 to 32 ranges are auto-configured for ATM channel groups and cannot be changed)

2.20.2.26 LAG Commands

lag

Syntax	[no] lag lag-id
Context	config
Description	This command creates the context for configuring Link Aggregation Group (LAG) attributes. A LAG can be used to group multiple ports into one logical link. The aggregation of multiple physical links allows for load sharing and offers seamless redundancy. If one of the links fails, traffic will be redistributed over the remaining links. Note that all ports in a LAG group must have autonegotiation set to Limited or Disabled. There are three possible settings for autonegotiation: • “on” or enabled with full port capabilities advertised • “off” or disabled where there is no autonegotiation advertisements • “limited” where a single speed/duplex is advertised. When autonegotiation is enabled on a port, the link attempts to automatically negotiate the link speed and duplex parameters. If autonegotiation is enabled, the configured duplex and speed parameters are ignored.

When autonegotiation is disabled on a port, the port does not attempt to autonegotiate and will only operate at the **speed** and **duplex** settings configured for the port. Note that disabling autonegotiation on gigabit ports is not allowed as the IEEE 802.3 specification for gigabit Ethernet requires autonegotiation be enabled for far end fault indication.

If the **autonegotiate limited** keyword option is specified the port will autonegotiate but will only advertise a specific speed and duplex. The speed and duplex advertised are the **speed** and **duplex** settings configured for the port. One use for limited mode is for multispeed gigabit ports to force gigabit operation while keeping autonegotiation is enabled for compliance with IEEE 801.3.

The system requires that autonegotiation be disabled or limited for ports in a LAG to guarantee a specific port speed.

LAG ID, ranging from 1 to 64, support up to 64 LAG members, and LAG ID above 64, support 32 LAG members.

The **no** form of this command deletes the LAG from the configuration. Deleting a LAG can only be performed while the LAG is administratively shut down. Any dependencies such as IP-Interfaces configurations must be removed from the configuration before issuing the **no lag** command.

Parameters *lag-id* — The LAG identifier, expressed as an integer.

Values 1 to 800

access

Syntax **access**

Context config>lag

Description This command enables the context to configure access parameters.

adapt-qos

Syntax **adapt-qos {link | port-fair | distribute [include-egr-hash-cfg]}**

Context config>lag>access

Description This command specifies how the LAG SAP queue and virtual scheduler buffering and rate parameters are adapted over multiple active XMA/MDAs. This command applies only to access LAGs.

Default adapt-qos distribute

Parameters **link** — Specifies that the LAG will create the SAP queues and virtual schedulers with the actual parameters on each LAG member port.

port-fair — Places the LAG instance into a mode that enforces QoS bandwidth constraints in the following manner:

- all egress QoS objects associated with the LAG instance are created on a per port basis
- bandwidth is distributed over these per port objects based on the proportion of the port's bandwidth relative to the total of all active ports bandwidth within the LAG
- the **include-egr-hash-cfg** behavior is automatically enabled allowing the system to detect objects that hash to a single egress link in the lag and enabling full bandwidth for that object on the appropriate port

distribute — Creates an additional internal virtual scheduler per IOM/XCM as parent of the configured SAP queues and virtual schedulers per LAG member port on that IOM/XCM. This internal virtual scheduler limits the total amount of egress bandwidth for all member ports on the IOM/XCM to the bandwidth specified in the egress qos policy.

include-egr-hash-cfg — Specifies whether explicitly configured hashing should factor into the egress buffering and rate distribution.

When this parameter is configured, all SAPs on this LAG which have explicit hashing configured, the egress HQos and HPol (including queues, policers, schedulers and arbiters) will receive 100% of the configured bandwidth (essentially operating in adapt-qos link mode). For any Multi-Service-Sites assigned to such a LAG, bandwidth will continue to be divided according to adapt-qos distribute mode.

A LAG instance that is currently in adapt-qos link mode may be placed at any time in port-fair mode. Similarly, a LAG instance that is currently in adapt-qos port-fair mode may be placed at any time in link mode. However, a LAG instance in adapt-qos distribute mode may not be placed into port-fair (or link) mode while QoS objects are associated with the LAG instance. To move from distribute to port-fair mode it is necessary to remove all QoS objects from the LAG instance.

bandwidth

Syntax	bandwidth <i>bandwidth</i> no bandwidth
Context	config>lag>access config>port>ethernet>access
Description	This command specifies the admin bandwidth assigned to SAPs, ports and LAGs which is used by SAP bandwidth CAC. The admin bandwidth on a port or LAG can be over or under booked using the booking-factor command.

Port or LAG: Increasing the port or LAG admin bandwidth will increase the available admin bandwidth on that port or LAG. Reducing the port or LAG admin bandwidth will reduce the available admin bandwidth on that port or LAG, however, if the reduction of available admin bandwidth would cause it to be insufficient to cover the sum of the current SAP admin bandwidth on the port or LAG then the command will fail.

The **no** version of the command reverts to the default value.

Default	no bandwidth
Parameters	<i>bandwidth</i> — Specifies the admin bandwidth assigned to the SAP, port or LAG, in kb/s.
Values	1 to 3200000000

booking-factor

Syntax	booking-factor <i>factor</i> no booking-factor
Context	config>lag>access config>port>ethernet>access
Description	This command specifies the booking factor applied against the port or LAG admin bandwidth by SAP admin bandwidth CAC. The service manager keeps track of the available admin bandwidth for each port or LAG configured with an admin bandwidth. The port or LAG available admin bandwidth is adjusted by the user configured booking factor, allowing the port or LAG bandwidth to be over or under booked. If the booking factor is increased then available admin bandwidth on the port or LAG increases. If the booking factor is decreased then available admin bandwidth on the port or LAG decreases, however, if the reduction of available admin bandwidth would cause it to be insufficient to cover the sum of the current SAP admin bandwidth on the port or LAG then the command fails. The no form of the command reverts to the default value.
Default	booking-factor 100
Parameters	<i>factor</i> — Specifies the percentage of the port or LAG admin bandwidth for SAP bandwidth CAC.
Values	1 to 1000

per-fp-egr-queuing

Syntax	[no] per-fp-egr-queuing
Context	config>lag>access
Description	This command specifies whether a more efficient method of queue allocation for LAG SAPs should be utilized. The no form of the command disables the method of queue allocation for LAG SAPs.

per-fp-ing-queuing

Syntax	[no] per-fp-ing-queuing
Context	config>lag>access
Description	This command specifies whether a more efficient method of queue allocation for LAG SAPs should be utilized. The no form of the command disables the method of queue allocation for LAG SAPs.

per-fp-sap-instance

Syntax	[no] per-fp-sap-instance
Context	config>lag>access
Description	This command enables optimized SAP instance allocation on a LAG. When enabled, SAP instance is allocated per each FP the LAG links exits on instead of per each LAG port. The no form of this command disables optimized SAP instance allocation.
Default	no per-fp-sap-instance

bfd

Syntax	bfd
Context	config>lag
Description	This command creates the bfd context and enables BFD over the associated LAG links.

disable-soft-reset-extension

Syntax	[no] disable-soft-reset-extension
Context	config>lag>bfd
Description	This command enables the BFD context and enables BFD over LAG links. Additional parameter configuration is required to make BFD over LAG links operational. Normally, BFD session timers are automatically extended during soft-reset operation on the IOMs and IMMs to avoid BFD sessions timing out and causing protocol events. However, in some cases this behavior is not desired as it could delay fast re-route transitions if they are in place. The optional disable-soft-reset-extension keyword allows this behavior to be disabled so that the BFD timers are not automatically extended.
Default	no disable-soft-reset-extension

Parameters	disable-soft-reset-extension — Disables the automatic extension of BFD timers during an IOM/IMM soft-reset.

family

Syntax	family [ipv4 ipv6]
Context	config>lag>bfd
Description	This command specifies the address family for the micro-BFD session over the associated LAG links.
Default	family ipv4
Parameters	ipv4 — Specifies that IPv4 encapsulation be used for the micro-BFD session. ipv6 — Specifies that IPv6 encapsulation be used for the micro-BFD session.

bfd-on-distributing-only

Syntax	[no] bfd-on-distributing-only
Context	config>lag>bfd>family
Description	This command enables restricting micro-BFD sessions to links in LACP state distributing. The no form of the command disables restricting micro-BFD sessions
Default	no bfd-on-distributing-only

local-ip-address

Syntax	local-ip-address ip-address no local-ip-address
Context	config>lag>bfd>family
Description	This command is used to specify the IPv4 or IPv6 address of the BFD source. The no form of the command removes this address from the configuration.
Default	no local-ip-address
Parameters	<i>ip-address</i> — Specifies the IP address. Values ipv4-address: a.b.c.d ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:d.d.d.d

x:-[0 to FFFF]H

d: [0 to 255]D

max-admin-down-time

Syntax	max-admin-down-time <i>[[down-interval] infinite]</i> no max-admin-down-time
Context	config>lag>bfd>family
Description	This command specifies the maximum amount of time the router will continue to forward traffic over a link after the micro-BFD sessions has transitioned to a Down state because it received an ADMIN-DOWN state from the far-end. This timer provide the administrator the configured amount of time to disable or de-provision the micro-BFD session on the local node before forwarding is halted over the associated link(s). The no form of the command removes the time interval from the configuration.
Default	max-admin-down-time 0
Parameters	<i>down-interval</i> — Specifies the amount of time, in seconds. Values -1 to 3600 infinite — Specifies no end time to forward traffic.

max-setup-time

Syntax	max-setup-time <i>[[up-interval] infinite]</i> no max-setup-time
Context	config>lag>bfd>family
Description	This command specifies the maximum amount of time the router will forward traffic over a link that has transitioned from Standby to Active, before the micro-BFD session must be fully established (Up state). The no form of the command returns the timer value to the default (0) which indicates that forwarding will not start until the BFD session is established.
Default	max-setup-time infinite
Parameters	<i>up-interval</i> — Specifies the amount of time, in milliseconds. Values -1 to 60000 infinite — Specifies no end time to forward traffic.

multiplier

Syntax	multiplier [<i>multiplier</i>] no multiplier
Context	config>lag>bfd>family
Description	This command specifies the detect multiplier used for a micro-BFD session over the associated LAG links. If a BFD control packet is not received for a period of multiplier X receive-interval then the session is declared down. The no form of the command removes multiplier from the configuration.
Default	multiplier 3
Parameters	<i>multiplier</i> — Specifies the multiplier value. Values 3 to 20

receive-interval

Syntax	receive-interval <i>receive-interval</i> no receive-interval
Context	config>lag>bfd>family
Description	This command specifies the receive timer used for micro-BFD session over the associated LAG links. The no form of the command removes the receive timer from the configuration.
Default	receive-interval 100
Parameters	<i>receive-interval</i> — Specifies the interval value, in milliseconds. Values 10 to 100000 Default 100 for CPM3 or later, 1000 for all others

remote-ip-address

Syntax	remote-ip-address <i>ip-address</i> no remote-ip-address
Context	config>lag>bfd>family
Description	This command specifies the IPv4 or IPv6 address of the BFD destination. The no form of the command removes this address from the configuration.

Default	no remote-ip-address
Parameters	<i>ip-address</i> — Specifies the IP address.
	Values
	ipv4-address: a.b.c.d
	ipv6-address: x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x:-[0 to FFFF]H
	d: [0 to 255]D

transmit-interval

Syntax	transmit-interval <i>transmit-interval</i> no transmit-interval
Context	config>lag>bfd>family
Description	This command specifies the transmit timer used for micro-BFD session over the associated LAG links. The no form of the command removes the transmit timer from the configuration.
Default	transmit-interval 100
Parameters	<i>transmit-interval</i> — Specifies the interval value, in milliseconds.
	Values 10 to 100000
	Default 100 for CPM3 or later, 1000 for all others

shutdown

Syntax	shutdown no shutdown
Context	config>lag>bfd>family
Description	This command disables micro BFD sessions for this address family. The no form of the command re-enables micro BFD sessions for this address family.
Default	shutdown

dynamic-cost

Syntax	[no] dynamic-cost

Context	config>lag
Description	This command enables OSPF or ISIS costing of a Link Aggregation Group (LAG) based on the available aggregated, operational bandwidth. The path cost is dynamically calculated based on the interface bandwidth. OSPF path cost can be changed through the interface metric or the reference bandwidth. If dynamic cost is configured, then costing is applied based on the total number of links configured and the cost advertised is inversely proportional to the number of links available at the time. This is provided that the number of links that are up exceeds the configured LAG threshold value at which time the configured threshold action determines if, and at what cost, this LAG will be advertised. For example: Assume a physical link in OSPF has a cost associated with it of 100, and the LAG consists of four physical links. The cost associated with the logical link is 25. If one link fails then the cost would automatically be adjusted to 33. If dynamic cost is not configured and OSPF autcost is configured, then costing is applied based on the total number of links configured. This cost will remain static provided the number of links that are up exceeds the configured LAG threshold value at which time the configured threshold action determines if and at what cost this LAG will be advertised. If dynamic-cost is configured and OSPF autcost is not configured, the cost is determined by the cost configured on the OSPF metric provided the number of links available exceeds the configured LAG threshold value at which time the configured threshold action determines if this LAG will be advertised. If neither dynamic-cost nor OSPF autcost are configured, the cost advertised is determined by the cost configured on the OSPF metric provided the number of links available exceeds the configured LAG threshold value at which time the configured threshold action determines if this LAG will be advertised. The no form of this command removes dynamic costing from the LAG.
Default	no dynamic-cost

encap-type

Syntax	encap-type {dot1q null qinq} no encap-type
Context	config>lag
Description	This command configures the encapsulation method used to distinguish customer traffic on a LAG. The encapsulation type is configurable on a LAG port. The LAG port and the port member encapsulation types must match when adding a port member.

If the encapsulation type of the LAG port is changed, the encapsulation type on all the port members will also change. The encapsulation type can be changed on the LAG port only if there is no interface associated with it. If the MTU is set to a non default value, it will be reset to the default value when the encap type is changed.

The **no** form of this command restores the default.

- Default** encap-type null — All traffic on the port belongs to a single service or VLAN.
- Parameters** **dot1q** — Ingress frames carry 802.1Q tags where each tag signifies a different service.
null — Ingress frames will not use any tags to delineate a service. As a result, only one service can be configured on a port with a null encapsulation type.
qinq — Specifies QinQ encapsulation.

hold-time

- Syntax** **hold-time down** *hold-down-time*
no hold-time
- Context** config>lag
- Description** This command specifies the timer, in tenths of seconds, which controls the delay between detecting that a LAG is down (all active ports are down) and reporting it to the higher levels.
- A non-zero value can be configured, for example, when active/standby signaling is used in a 1:1 fashion to avoid informing higher levels during the small time interval between detecting that the LAG is down and the time needed to activate the standby link.
- Default** no hold-time
- Parameters** *hold-down-time* — Specifies the hold-time for event reporting.
- Values** 0 to 2000

lACP

- Syntax** **lACP** [**mode**] [**administrative-key** *admin-key*] [**system-id** *system-id*] [**system-priority** *priority*]
- Context** config>lag
- Description** This command specifies the LACP mode for aggregated Ethernet interfaces only. This command enables the LACP protocol. Per the IEEE 802.1ax standard, the Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information between Partner Systems on a link to allow their Link Aggregation Control instances to reach agreement on the identity of the Link Aggregation Group to which the link belongs, move the link to that Link Aggregation Group, and enable its transmission and reception functions in an orderly manner.

Note that if any of the parameters are omitted, the existing configuration is preserved. The default parameter values are used if a parameter is never explicitly configured.

Default	no lacp
Parameters	mode — Specifies the mode in which LACP will operate. Values passive — Starts transmitting LACP packets only after receiving packets. active — Initiates the transmission of LACP packets. <i>admin-key</i> — Specifies an administrative key value to identify the channel group on each port configured to use LACP. This value should be configured only in exceptional cases. A random key is assigned by default if a value is not specified. Values 1 to 65535 <i>system-id</i> — Specifies the 48-bit system ID in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses. Values 1 to 65535 Default 32768 <i>priority</i> — Specifies the system priority. Values 1 to 65535 Default 32768

lacp-mux-control

Syntax	lacp-mux-control {coupled independent} no lacp-mux-control
Context	config>lag
Description	This command configures the type of multiplexing machine control to be used in a LAG with LACP in active/passive modes. The no form of the command disables multiplexing machine control.
Default	lacp-mux-control coupled
Parameters	coupled — TX and RX activate together. independent — RX activates independent of TX.

lacp-xmit-interval

Syntax	lacp-xmit-interval {slow fast}

no lacp-xmit-interval

Context	config>lag
Description	This command specifies the interval signaled to the peer and tells the peer at which rate it should transmit.
Default	lacp-xmit-interval fast
Parameters	slow — Transmits packets every 30 seconds. fast — Transmits packets every second.

lacp-xmit-stdby

Syntax	[no] lacp-xmit-stdby
Context	config>lag
Description	This command enables LACP message transmission on standby links. The no form of this command disables LACP message transmission. This command should be disabled for compatibility when using active/standby groups. This forces a timeout of the standby links by the peer. Use the no form if the peer does not implement the correct behavior regarding the lacp sync bit.
Default	lacp-xmit-stdby

link-map-profile

Syntax	link-map-profile <i>link-map-profile-id</i> [create] no link-map-profile <i>link-map-profile-id</i>
Context	config>lag
Description	This command creates the link map profile that can to control which LAG ports are to be used on egress or enables the configuration context for previously created link map profile. link map profiles are not created by default. The no form of this command, deletes the specified link map profile.
Parameters	<i>link-map-profile-id</i> — An integer from 1 to 64 that defines a unique LAG link map profile on this LAG.

link

Syntax	link <i>port-id</i> { primary secondary }

	no link <i>port-id</i>																								
Context	config>lag>link-map-profile																								
Description	This command designates one of the configured ports of the LAG to be used on egress as either a primary or secondary link (based on the option selected) by all SAPs and network interfaces that use this LAG link map profile. Links are part of a profile. When a link is added or deleted, all SAPs and network interfaces that use this link-map-profile may be re-hashed if required. The no form of this command deletes the link from this LAG link mapping profile. A port must be deleted from all LAG link profiles if it is to be deleted from the LAG.																								
Parameters	<i>port-id</i> — Specifies a physical port ID that is an existing member of this LAG. <table><tr><td><i>port-id</i></td><td><i>slot/mda/port[.channel]</i></td><td></td></tr><tr><td>eth-sat-id</td><td>esat-id/slot/port</td><td></td></tr><tr><td></td><td>esat</td><td>keyword</td></tr><tr><td></td><td>id</td><td>1 to 20</td></tr><tr><td>pxc-id</td><td>pxc-id.sub-port</td><td></td></tr><tr><td></td><td>pxc</td><td>keyword</td></tr><tr><td></td><td>id</td><td>1 to 64</td></tr><tr><td></td><td>sub-port</td><td>a, b</td></tr></table>	<i>port-id</i>	<i>slot/mda/port[.channel]</i>		eth-sat-id	esat-id/slot/port			esat	keyword		id	1 to 20	pxc-id	pxc-id.sub-port			pxc	keyword		id	1 to 64		sub-port	a, b
<i>port-id</i>	<i>slot/mda/port[.channel]</i>																								
eth-sat-id	esat-id/slot/port																								
	esat	keyword																							
	id	1 to 20																							
pxc-id	pxc-id.sub-port																								
	pxc	keyword																							
	id	1 to 64																							
	sub-port	a, b																							
	primary — Designates one of the configured ports of the LAG to be used on egress as a primary link by SAPs/network interfaces that use this LAG link map profile. secondary — Designates one of the configured ports of the LAG to be used on egress as a secondary link by SAPs/network interfaces that use this LAG link map profile.																								

failure-mode

Syntax	failure-mode [discard per-link-hash] no failure-mode
Context	config>lag>link-map-profile
Description	This command defines the failure mode for egress traffic of SAPs/network interfaces that use this link-map-profile when neither primary nor secondary links of this profile are available.
Default	failure-mode per-link-hash
Parameters	discard — Specifies egress traffic for SAPs/network interfaces using this link-map-profile is discarded to protect SAP/network interface traffic on other LAG links from impact of re-hashing the affected SAPs/network interfaces.

per-link-hash — Specifies egress traffic for SAPs/network interfaces using this link-map-profile is rehashed on remaining, available LAG links using per-link-hash algorithm. SAP/network interface QoS configurations dictate what traffic is discarded on any link that may become oversubscribed as result of the re-hash.

port

Syntax	port <i>port-id</i> [<i>port-id</i>] [priority <i>priority</i>] [sub-group <i>sub-group-id</i>] no port <i>port-id</i> [<i>port-id</i>]
Context	config>lag
Description	This command adds ports to a Link Aggregation Group (LAG). The port configuration of the first port added to the LAG is used as a basis to compare to subsequently added ports. If a discrepancy is found with a newly added port, that port will not be added to the LAG. Multiple (space separated) ports can be added or removed from the LAG link assuming the maximum of number of ports is not exceeded. Ports that are part of a LAG must be configured with auto-negotiate limited or disabled. The no form of this command removes ports from the LAG.
Default	No ports are defined as members of a LAG.
Parameters	<i>port-id</i> — The port ID configured or displayed. Note that the maximum number of ports in a LAG depends on platform-type, hardware deployed, and SR OS software release. Adding a port over the maximum allowed per given router/switch is blocked. Some platforms support double port scale for some port types on LAGs with lag-id in the range of 1 to 64 inclusive. Up to 16 ports can be specified in a single statement, up to 64 ports total. Values These values apply to the 7950 XRS only.

<i>slot/mda/port</i>		
eth-sat-id	<i>esat-id/slot/port</i>	
	esat	keyword
	id	1 to 20
pxc-id	<i>pxc-id.sub-port</i>	
	pxc	keyword
	id	1 to 64
	sub-port	a to b

priority — Port priority used by LACP. The port priority is also used to determine the primary port. The port with the lowest priority is the primary port. In the event of a tie, the smallest port ID becomes the primary port.

Values 1 to 65535

sub-group-id — This parameter identifies a LAG subgroup. When using subgroups in a LAG, they should only be configured on one side of the LAG, not both. Only having one side perform the active/standby selection will guarantee a consistent selection and fast convergence. The active/standby selection will be signaled through LACP to the other side. The hold time should be configured when using subgroups to prevent the LAG going down when switching between active and standby subgroup since momentarily all ports are down in a LAG (break-before-make).

Values 1 to 8 identifies a LAG subgroup
The **auto-iom** subgroup is defined based on the IOM (all ports of the same IOM are assigned to the same subgroup).
The **auto-md** subgroup is defined based on the MDA. (all ports of the same MDA are assigned to the same subgroup).

per-link-hash

Syntax	per-link-hash per-link-hash weighted [auto-rebalance] no per-link-hash
Context	config>lag
Description	This command configures per-link-hashing on a LAG. When enabled, SAPs/subscribers/interfaces are hashed on LAG egress to a single LAG link. The no form of this command disables per-link-hashing on a LAG.
Default	no per-link-hash
Parameters	weighted — SAPs/subscribers/interfaces are distributed amongst LAG links based on their preconfigured class and weight. As new links are added to a LAG, existing SAPs/subscribers/interfaces are not impacted. weighted auto-rebalance — SAPs/subscribers/interfaces are distributed amongst LAG links based on their preconfigured class and weight. As new links are added to a LAG, existing SAPs/subscribers/interfaces are rebalanced automatically.

port-threshold

Syntax	port-threshold <i>value</i> [action { dynamic-cost static-cost down }] [cost <i>static-cost</i>] no port-threshold
Context	config>lag

Description	This command configures the behavior for the Link Aggregation Group (LAG) if the number of operational links is equal to or below a threshold level. Nokia recommends that operators use weight-threshold instead of port-threshold to control LAG operational status. For example, when 10GE and 100GE ports are mixed in a LAG, each 10GE port will have a weight of 1, while each 100GE port will have a weight of 10. The weight-threshold command can also be used for LAGs with all ports of equal speed to allow a common operational model. For example, each port has a weight of 1 to mimic port-threshold and its related configuration. The no form of this command reverts to the default values.
Default	port-threshold 0 action down
Parameters	value — The decimal integer threshold number of operational links for the LAG at or below which the configured action will be invoked. If the number of operational links exceeds the port-threshold value, any action taken for being below the threshold value will cease. Values 0 to 63 action — Specifies the action to take if the number of active links in the LAG is at or below the threshold value. dynamic-cost — Specifies that dynamic costing will be activated. As a result, the LAG will remain operationally up with a cost relative to the number of operational links. The link will only be regarded as operationally down when all links in the LAG are down. static-cost — Specifies that static costing will be activated. As a result, the LAG will remain operationally up with the configured cost, regardless of the number of operational links. The link will only be regarded as operationally down when all links in the LAG are down. If this parameter is used with an IGP, its reference-bandwidth must also be configured. down — Specifies the LAG will be brought operationally down if the number of operational links is equal to or less than the configured threshold value. The LAG will only be regarded as up once the number of operational links exceeds the configured threshold value. static-cost — The decimal integer static cost of the LAG. Values 1 to 16777215

port-type

Syntax	port-type <i>lag-port-type</i> no port-type
Context	config>lag

Description	This command configures the port type for the link aggregation group. The no form of the command reverts to the default.
Default	port-type standard
Parameters	<i>lag-port-type</i> — Specifies the type of ports allowed in this LAG.
Values	standard — Allows all non-HSMDA and non-HS type ports to be added to this LAG hsmda — Limits the LAG members to be HSMDA-v2 ports hs — Limits the LAG members to be HSQ IOMs (iom4-e-hs) ports

port-weight-speed

Syntax	port-weight-speed <i>gbps</i> no port-weight-speed
Context	config>lag
Description	This command enables mixed port-speed LAG operation. Parameter specified with the command defines what type of ports are allowed in a LAG, and what is the weight of each port for total LAG weight calculation. The no form specifies that all LAG links must be of the same speed. Each link weight is 1. The no form disables mixed port-speed LAG operation if there are no mixed-speed links. Issuing this command automatically checks that all links are the same speed and re-calibrates the link weights. If all links are not the same speed, no-port-weight-speed is rejected.
Default	no port-weight-speed
Parameters	<i>gbps</i> — The port weight speed in Gigabits per second. Values 1, 10 port-weight-speed 1 — LAG supports any mix of 1GE, 10GE ports up to a total weight of 64 (for 64 link LAGs) or 32 (for 32 link LAGs). Each 1 GE port has a weight of 1; each 10GE port has a weight of 10. port-weight-speed 10 — LAG supports any mix of 10GE, 40GE, 100GE ports up to a total weight of 64 (for 64 link LAGs) or 32 (for 32 link LAGs). Each 10 GE port has a weight of 1; each 40GE port has a weight of 4; each 100GE port has a weight of 10. For existing LAGs: no port-weight-speed can be changed to port-weight-speed 1 when the LAG consists of only 1GE links. no port-weight-speed can be changed to port-weight-speed 10 when the LAG consists of only 10GE links. port-weight-speed 1 or port-weight-speed 10 can be changed to no port-weight-speed in service, when all links of the LAG are 1GE, 10GE, 40GE, or 100GE.

All other configuration changes require shutdown of the LAG and removal of all ports first.

selection-criteria

Syntax	selection-criteria [best-port highest-count highest-weight] [slave-to-partner] [subgroup-hold-time <i>hold-time</i>] no selection-criteria						
Context	config>lag						
Description	This command specifies which selection criteria should be used to select the active sub-group. If there is a tie for highest-count or highest-weight, the LAG will prefer the port with the lowest priority. If that does not break the tie, the currently active sub group will stay active (that is, non-revertive behavior). The no form of this command reverts to the default value.						
Default	selection-criteria highest-count						
Parameters	highest-count — Selects a sub-group with the highest number of eligible members as an active sub-group (not applicable to “power-off” mode of operations). highest-weight — Selects a sub-group with the highest aggregate weight as an active subgroup (not applicable to “power-off” mode of operations). Aggregate weight is calculated as the sum of (65535 - port priority) all ports within a sub-group. best-port — Selects a sub-group containing the port with highest priority port as an active subgroup. In case of equal port priorities, the sub-group containing the port with the lowest port-id is chosen. slave-to-partner — The slave-to-partner keyword specifies that it, together with the selection criteria, should be used to select the active sub-group. An eligible member is a LAG-member link which can potentially become active. This means it is operationally up (not disabled) for use by the remote side. The slave-to-partner parameter can be used to control whether or not this latter condition is taken into account. <i>hold-time</i> — Applicable with LACP enabled. Specifies the optional delay timer for switching to a newly selected active sub-group from the existing active sub-group. The timer delay applies only if the existing sub-group remains operationally up. Values <table> <tr> <td>not specified</td><td>Equivalent to specifying a value of 0. Specifies no delay and to switchover immediately to a new candidate active sub-group.</td></tr> <tr> <td>0 to 2000</td><td>Integer specifying the timer value in 10ths of a second.</td></tr> <tr> <td>infinite</td><td>Do not switchover from existing active sub-group if the subgroup remains UP. Manual switchover possible using tools perform lag force command.</td></tr> </table>	not specified	Equivalent to specifying a value of 0. Specifies no delay and to switchover immediately to a new candidate active sub-group.	0 to 2000	Integer specifying the timer value in 10ths of a second.	infinite	Do not switchover from existing active sub-group if the subgroup remains UP. Manual switchover possible using tools perform lag force command.
not specified	Equivalent to specifying a value of 0. Specifies no delay and to switchover immediately to a new candidate active sub-group.						
0 to 2000	Integer specifying the timer value in 10ths of a second.						
infinite	Do not switchover from existing active sub-group if the subgroup remains UP. Manual switchover possible using tools perform lag force command.						

standby-signaling

Syntax	standby-signaling {lacp power-off} no standby-signaling
Context	config>lag
Description	This command specifies how the state of a member port is signaled to the remote side when the status corresponding to this member port has the standby value.
Default	standby-signaling lacp

weight-threshold

Syntax	weight-threshold <i>value</i> [action {dynamic-cost static-cost down}] [cost <i>static-cost</i>] no weight-threshold
Context	config>lag
Description	This command configures the behavior for the Link Aggregation Group (LAG) if the total weight of operational links is equal to or below the configured threshold level. The command can be used for mixed port-speed LAGs and for LAGs with all ports of equal speed. The no form of this command disabled weight-threshold operation in LAG.
Default	weight-threshold 0 action down
Parameters	<i>value</i> — The decimal integer threshold number of operational links for the LAG at or below which the configured action will be invoked. If the number of operational links exceeds the weight-threshold value, any action taken for being below the threshold value will cease. Values 0 to 63 action — Specifies the action to take if the number of active links in the LAG is at or below the threshold value. dynamic-cost — Specifies that dynamic costing will be activated. As a result the LAG will remain operationally up with a cost relative to the number of operational links. The link will only be regarded as operationally down when all links in the LAG are down. static-cost — Specifies that static costing will be activated. As a result, the LAG will remain operationally up with the configured cost, regardless of the number of operational links. The link will only be regarded as operationally down when all links in the LAG are down. If this parameter is used with an IGP, its reference-bandwidth must also be configured.

down — Specifies the LAG will be brought operationally down if the number of operational links is equal to or less than the configured threshold value. The LAG will only be regarded as up once the number of operational links exceeds the configured threshold value.

static-cost — The decimal integer static cost of the LAG.

Values 1 to 16777215

2.20.2.27 Eth Tunnel Commands

eth-tunnel

Syntax `[no] eth-tunnel tunnel-index`

Context config

Description This command configures a G.8031 protected Ethernet tunnel.

The **no** form of this command deletes the Ethernet tunnel specified by the tunnel-id.

Parameters *tunnel-index* — Specifies the tunnel index.

Values 1 to 1024

ccm-hold-time

Syntax `ccm-hold-time {down down-timeout | up up-timeout}`
`no ccm-hold-time`

Context config>eth-tunnel

Description This command configures eth-tunnel CCM dampening timers.

The no form of the command reverts to the default.

Default no ccm-hold-time

Parameters *down-timeout* — Specifies the eth-tunnel CCM down timers, in 100ths of seconds.

Values 0 to 1000

Default 0

up-timeout — Specifies the eth-tunnel CCM up timers, in 10ths of seconds.

Values 0 to 5000

Default 20

description

Syntax	description <i>long-description-string</i> no description
Context	config>eth-tunnel
Description	This command adds a text description for the eth-tunnel. The no form of this command removes the text description.
Default	Eth-tunnel
Parameters	<i>long-description-string</i> — Specifies the text description up to 160 characters.

ethernet

Syntax	ethernet
Context	config>eth-tunnel
Description	This command enables the context to configure Ethernet parameters.

encap-type

Syntax	encap-type {dot1q qinq} no encap-type
Context	config>eth-tunnel>ethernet
Description	This command configures the encapsulation method.
Default	encap-type dot1q
Parameters	dot1q — Specifies dot1q encapsulation. qinq — Specifies qinq encapsulation.

mac

Syntax	mac <i>ieee-address</i> no mac
Context	config>eth-tunnel>ethernet

Description	This command assigns a specific MAC address to an Ethernet port, Link Aggregation Group (LAG), Ethernet tunnel or BCP-enabled port or sub-port. Only one MAC address can be assigned to a port. When multiple mac commands are entered, the last command overwrites the previous command. When the command is issued while the port is operational, IP will issue an ARP, if appropriate, and BPDUs are sent with the new MAC address. The no form of this command returns the MAC address to the default value.
Default	A default MAC address is assigned by the system from the chassis MAC address pool.
Parameters	<i>ieee-address</i> — Specifies the source MAC address.

lag-emulation

Syntax	lag-emulation
Context	config>eth-tunnel
Description	This command configures eth-tunnel load sharing parameters.

access

Syntax	access
Context	config>eth-tunnel>lag-emulation
Description	This command configures eth-tunnel load sharing access parameters

adapt-qos

Syntax	adapt-qos {distribute link port-fair} no adapt-qos
Context	config>eth-tunnel>lag-emulation>access
Description	This command configures how the Ethernet Tunnel group SAP queue and virtual scheduler buffering and rate parameters are adapted over multiple active MDAs. The no form of the command reverts the default.
Default	no adapt-qos
Parameters	distribute — Each MDA will receive a fraction of the SAP and scheduler parameters. link — The Ethernet Tunnel group will create the SAP queues and virtual schedulers with the actual parameters on each MDA.

port-fair — Places the LAG instance into a mode that enforces QoS bandwidth constraints in the following manner:

- All egress QoS objects associated with the LAG instance are created on a per port basis.
- Bandwidth is distributed over these per port objects based on the proportion of the port's bandwidth relative to the total of all active ports bandwidth within the LAG.
- The inc-egr-hash-cfg behavior is automatically enabled allowing the system to detect objects that hash to a single egress link in the lag and enabling full bandwidth for that object on the appropriate port.

A LAG instance that is currently in adapt-qos link mode may be placed at any time in port-fair mode. Similarly, a LAG instance currently in adapt-qos port-fair mode may be placed at any time in link mode. However, a LAG instance in adapt-qos distribute mode may not be placed into port-fair (or link) mode while QoS objects are associated with the LAG instance. To move from distribute to port-fair mode either remove all QoS objects from the LAG instance or remove all member ports from the LAG instance.

per-fp-ing-queuing

Syntax	[no] per-fp-ing-queuing
Context	config>eth-tunnel>lag-emulation>access
Description	This command configures whether a more efficient method of queue allocation for Ethernet Tunnel Group SAPs should be utilized. The no form of the command reverts the default.
Default	no per-fp-ing-queuing

path-threshold

Syntax	path-threshold <i>num-paths</i> no path-threshold
Context	config>eth-tunnel>lag-emulation
Description	This command configures the behavior for the eth-tunnel if the number of operational members is equal to or below a threshold level
Default	no path-threshold
Parameters	<i>num-paths</i> — Specifies the threshold for the Ethernet Tunnel group. Values 0 to 15

protection-type

Syntax	protection-type { g8031-1to1 loadsharing }
Context	config>eth-tunnel
Description	This command configures the model used for determining which members are actively receiving and transmitting data. The no form of the command reverts to the default.
Default	protection-type g8031-1to1
Parameters	g8031-1to1 — As per G.8031 spec, only two members are allowed, and only one of them can be active at one point in time. loadsharing — Multiple members can be active at one point in time.

revert-time

Syntax	revert-time <i>time</i> no revert-time
Context	config>eth-tunnel
Description	This command configure how long to wait before switching back to the primary path after it has been restored to Ethernet tunnel. The no form of this command disables the revert behavior, effectively setting the revert time to zero.
Default	no revert-time
Parameters	<i>time</i> — Specifies the re-activation delay, in seconds, for the primary path. Values 1 to 720

path

Syntax	[no] path <i>path-index</i>
Context	config>eth-tunnel
Description	This command configures one of the two paths supported under the Ethernet tunnel. Although the values indicate 1 to 8, only two paths, 1 and 2, are currently supported. The no form of this command removes the path from under the Ethernet tunnel. If this is the last path, the associated SAP needs to be un-configured before the path can be deleted.
Default	no path

Parameters *path-index* — Specifies the identifier for the path.
Values 1 to 16

description

Syntax **description** *description-string*
no description

Context config>eth-tunnel>path

Description This command configures a text description for the path.
The **no** form of this command removes the text description.

Default no description

Parameters *description-string* — Specifies a text description.
Values Maximum 80 characters

member

Syntax **member** *port-id*
no member

Context config>eth-tunnel>path

Description This command associates a port with the path defined under the Ethernet tunnel. If the operator wants to replace an existing member port or control tag, the whole path needs to be shutdown first. The alternate path will be activated as a result keeping traffic interruption to a minimum. Then the whole path must be deleted, the alternate path precedence modified to primary before re-creating the new path.

The following port-level configuration needs to be the same across the two member ports of an Ethernet tunnel:

- port>ethernet>access>{ingress|egress}>queue-group
- port>ethernet>egress-scheduler-policy
- port>access>egress>pool
- port>ethernet>dot1q-etype
- port>ethernet>qinq-etype
- port>ethernet>pbb-etype
- port>ethernet>mtu

The Ethernet tunnel will inherit the configuration from the first member port for these parameters. Additional member port that is added must have the same configuration.

The operator is allowed to update these port parameters only if the port is the sole member of an Ethernet tunnel. This means that in the example below, the operator needs to remove port 1/1/4 and port 1/1/5 before being allowed to modify 1/1/1 for the above parameters.

```
eth-tunnel 1
  path 1
    member 1/1/1
  path 2
    member 1/1/4
eth-tunnel 2
  path 1
    member 1/1/1
  path 2
    member 1/1/5
```

The **no** form of this command is used just to indicate that a member is not configured. The procedure described above, based on the **no path** command must be used to un-configure/change the member port assigned to the path.

Default no member

Parameters *port-id* — Specifies the port-id associated with the path in the format x/y/z where x represents the IOM, y the MDA and z the port numbers.

control-tag

Syntax **control-tag** *qtag* [*qtag*]
no control-tag

Context config>eth-tunnel>path

Description This command specifies the VLAN-ID to be used for Ethernet CFM and G.8031 control plane exchanges. If the operator wants to replace an existing control-tag, the parent path needs to be in shutdown state, then deleted and recreated before a new control-tag can be specified.

The **no** form of this command is used just to indicate that a control-tag is not configured. The procedure described above, based on 'no path' command must be used to un-configure/change the control-tag assigned to the path.

Default no control-tag

Parameters *qtag* [*qtag*] — Specifies the value of the VLAN ID to be used for the control tag.

Values 1 to 4094, untagged option is not supported, *

precedence

Syntax **precedence** {primary | secondary}
no precedence

Context	config>eth-tunnel>path
Description	This command specifies the precedence to be used for the path. Only two precedence options are supported: primary and secondary . The no form of this command sets the precedence to the default value.
Default	precedence secondary
Parameters	primary secondary — Specifies the path precedence as either primary or secondary.

eth-cfm

Syntax	eth-cfm
Context	config>eth-tunnel>path
Description	This command enables the context to configure ETH-CFM parameters.

mep

Syntax	[no] mep mep-id domain md-index association ma-index
Context	config>eth-tunnel>path>eth-cfm
Description	This command provisions an 802.1ag maintenance endpoint (MEP). The no form of the command reverts to the default values.
Parameters	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 to 8191 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 to 4294967295 <i>ma-index</i> — Specifies the MA index value. Values 1 to 4294967295

control-mep

Syntax	[no] control-mep
Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command enables the Ethernet ring control on the MEP. The use of control-mep command is mandatory for a ring. MEP detection of failure using CCM may be enabled or disabled independently of the control mep.

The **no** form of this command disables Ethernet ring control.

ccm-enable

Syntax	[no] ccm-enable
Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command enables the generation of CCM messages. The no form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax	ccm-ltm-priority <i>priority</i> no ccm-ltm-priority
Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command specifies the priority value for CCMs and LTMs transmitted by the MEP. The no form of the command removes the priority value from the configuration.
Default	The highest priority on the bridge-port.
Parameters	<i>priority</i> — Specifies the priority of CCM and LTM messages. Values 0 to 7

eth-test-enable

Syntax	[no] eth-test-enable
Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command enables eth-test functionality on MEP. For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands: oam eth-cfm eth-test <i>mac-address mep mep-id domain md-index association ma-index</i> [priority <i>priority</i>] [data-length <i>data-length</i>] A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax	test-pattern { all-zeros all-ones } [crc-enable] no test-pattern
Context	config>eth-tunnel>path>eth-cfm>mep>eth-test-enable
Description	This command configures the test pattern for eth-test frames. The no form of the command removes the values from the configuration.
Parameters	all-zeros — Specifies to use all zeros in the test pattern. all-ones — Specifies to use all ones in the test pattern. crc-enable — Generates a CRC checksum. Default all-zeros

low-priority-defect

Syntax	low-priority-defect { allDef macRemErrXcon remErrXcon errXcon xcon noXcon }												
Context	config>eth-tunnel>path>eth-cfm>mep												
Description	This command specifies the lowest priority defect that is allowed to generate a fault alarm.												
Default	low-priority-defect remErrXcon												
Parameters	allDef macRemErrXcon remErrXcon errXcon xcon noXcon — Specifies the lowest priority defect.												
	Values												
	<table> <tr> <td>allDef</td><td>DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM</td></tr> <tr> <td>macRemErrXcon</td><td>Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM</td></tr> <tr> <td>remErrXcon</td><td>Only DefRemoteCCM, DefErrorCCM, and DefXconCCM</td></tr> <tr> <td>errXcon</td><td>Only DefErrorCCM and DefXconCCM</td></tr> <tr> <td>xcon</td><td>Only DefXconCCM; or</td></tr> <tr> <td>noXcon</td><td>No defects DefXcon or lower are to be reported</td></tr> </table>	allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM	macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM	remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM	errXcon	Only DefErrorCCM and DefXconCCM	xcon	Only DefXconCCM; or	noXcon	No defects DefXcon or lower are to be reported
allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM												
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM												
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM												
errXcon	Only DefErrorCCM and DefXconCCM												
xcon	Only DefXconCCM; or												
noXcon	No defects DefXcon or lower are to be reported												

mac-address

Syntax	mac-address <i>mac-address</i>

no mac-address

Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command specifies the MAC address of the MEP. The no form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke SDP).
Parameters	<i>mac-address</i> — Specifies the MAC address of the MEP.
Values	6-byte unicast mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) of the MEP. Using the all zeros address is equivalent to the no form of this command.

control-mep

Syntax	[no] control-mep
Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command enables the usage of the CC state by the Ethernet tunnel manager for consideration in the protection algorithm. The use of control-mep command is recommended if fast failure detection is required, especially when Link Layer OAM does not provide the required detection time. The no form of this command disables the use of the CC state by the Ethernet tunnel manager.
Default	no control-mep

shutdown

Syntax	[no] shutdown
Context	config>eth-tunnel>path>eth-cfm>mep
Description	This command administratively enables/disables the MEP. The no form of this command enables the MEP.
Default	shutdown

shutdown

Syntax	[no] shutdown
Context	config>eth-tunnel>path

	config>eth-tunnel
Description	This command administratively enables/disables the path. The no form of this command enables the path.
Default	shutdown

2.20.2.28 ETH-CFM Configuration Commands

eth-cfm

Syntax	eth-cfm
Context	config>port>ethernet config>lag
Description	This command enables the context to configure 802.1ag CFM parameters.

mep

Syntax	[no] mep mep-id domain md-index association ma-index [vlan vlan-id]
Context	config>port>ethernet>eth-cfm config>lag>eth-cfm config>router>if>eth-cfm
Description	This command provisions the maintenance endpoint (MEP). The no form of the command reverts to the default values.
Parameters	<i>mep-id</i> — Specifies the maintenance association end point identifier. Values 1 to 8191 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 to 4294967295 <i>ma-index</i> — Specifies the MA index value. Values 1 to 4294967295 <i>vlan-id</i> — Specific to tunnel facility MEPs which means this option is only applicable to the lag>eth-cfm> context. Used to specify the outer vlan id of the tunnel. Values 1 to 4094

ais-enable

Syntax	[no] ais-enable
Context	config>port>ethernet>eth-cfm>mep config>lag>eth-cfm>mep
Description	This command enables the reception of AIS messages. The no form of the command reverts to the default values.

client-meg-level

Syntax	client-meg-level <i>[[level/ [level/]]</i> no client-meg-level				
Context	config>port>ethernet>eth-cfm>mep>ais-enable config>lag>eth-cfm> mep>ais-enable				
Description	This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level. Only the lowest client MEG level will be used for facility MEPs. The no form of the command reverts to the default values.				
Parameters	<i>level</i> — Specifies the client MEG level. <table><tr><td>Values</td><td>1 to 7</td></tr><tr><td>Default</td><td>1</td></tr></table>	Values	1 to 7	Default	1
Values	1 to 7				
Default	1				

interface-support-enable

Syntax	[no] interface-support-enable
Context	config>port>ethernet>eth-cfm>mep>ais-enable
Description	This command enables and disables the generation of AIS PDUs based on the associated endpoint state.

interval

Syntax	interval {1 60} no interval
Context	config>port>ethernet>eth-cfm>mep>ais-enable

config>lag>eth-cfm> mep>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

The **no** form of the command reverts to the default values.

Parameters **1 | 60** — The transmission interval of AIS messages, in seconds.

Default 1

low-priority-defect

Syntax **low-priority-defect {allDef | macRemErrXcon}**

Context config>lag>eth-cfm>mep>ais-enable

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default low-priority-defect remErrXcon

Parameters **allDef | macRemErrXcon** — Specifies the lowest priority defect.

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM
macRemErrXcon	Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
errXcon	Only DefErrorCCM and DefXconCCM
xcon	Only DefXconCCM; or
noXcon	No defects DefXcon or lower are to be reported

priority

Syntax **priority *priority-value***
no priority

Context config>port>ethernet>eth-cfm>mep>ais-enable
config>lag>eth-cfm> mep>ais-enable

Description This command specifies the priority of the AIS messages generated by the node.

The **no** form of the command reverts to the default values.

Parameters	<i>priority-value</i> — Specifies the priority value of the AIS messages originated by the node.
Values	0 to 7
Default	7

alarm-notification

Syntax	alarm-notification
Context	config>eth-tunnel>path>eth-cfm>mep config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep
Description	This command configures the MEP alarm notification parameter.

fng-alarm-time

Syntax	fng-alarm-time <i>time</i>
Context	config>eth-tunnel>path>eth-cfm>mep>alarm-notification config>lag>eth-cfm>mep>alarm-notification config>port>ethernet>eth-cfm>mep>alarm-notification
Description	This command configures the Fault Notification Generation (FNG) alarm time.
Parameters	<i>time</i> — The length of time, in centi-seconds, that must expire before a defect is alarmed.
Values	0, 250, 500, 1000
Default	0

fng-reset-time

Syntax	fng-reset-time <i>time</i>
Context	config>eth-tunnel>path>eth-cfm>mep>alarm-notification config>lag>eth-cfm>mep>alarm-notification config>port>ethernet>eth-cfm>mep>alarm-notification
Description	This command configure the Fault Notification Generation (FNG) reset time.
Parameters	<i>time</i> — The length of time, in centiseconds, that must expire before a defect is reset.
Values	0, 250, 500, 1000
Default	0

ccm-enable

Syntax	[no] ccm-enable
Context	config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep
Description	This command enables the generation of CCM messages. The no form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax	ccm-ltm-priority <i>priority</i> no ccm-ltm-priority				
Context	config>lag>eth-cfm>mep> config>port>ethernet>eth-cfm>mep> config>router>if>eth-cfm>mep				
Description	This command specifies the priority of the CCM and LTM messages transmitted by the MEP. Since CCM does not apply to the Router Facility MEP only the LTM priority is of value under that context. The no form of the command reverts to the default values.				
Default	no ccm-ltm-priority				
Parameters	<i>priority</i> — Specifies the priority value. <table> <tr> <td>Values</td><td>0 to 7</td></tr> <tr> <td>Default</td><td>7, highest priority for CCMs and LTMs transmitted by the MEP</td></tr> </table>	Values	0 to 7	Default	7, highest priority for CCMs and LTMs transmitted by the MEP
Values	0 to 7				
Default	7, highest priority for CCMs and LTMs transmitted by the MEP				

ccm-padding-size

Syntax	ccm-padding-size <i>ccm-padding</i> no ccm-padding-size		
Context	config>eth-tunnel>path>eth-cfm>mep config>lag>eth-cfm>mep		
Description	This command inserts additional padding in the CCM packets. The no form of the command reverts to the default.		
Parameters	<i>ccm-padding</i> — Specifies the additional padding in the CCM packets, in octets. <table> <tr> <td>Values</td><td>3 to 1500</td></tr> </table>	Values	3 to 1500
Values	3 to 1500		

ccm-tlv-ignore

Syntax	ccm-tlv-ignore [port-status] [interface-status] no ccm-tlv-ignore
Context	config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep
Description	This command allows the receiving MEP to ignore the specified TLVs in CCM PDU. Ignored TLVs will be reported as absent and will have no impact on the MEP state machine. The no form of the command causes the receiving MEP will process all recognized TLVs in the CCM PDU.
Parameters	port-status — Ignores the port status TLV on reception. interface-status — Ignores the interface status TLV on reception.

collect-imm-stats

Syntax	[no] collect-imm-stats
Context	config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep config>router>if>eth-cfm>mep
Description	This command enables the collection of statistics on the facility MEPs. This command is an object under the Facility MEP. This is at a different level of the hierarchy than collection of Imm statistics for service SAPs and MPLS SDP Bindings. The show mep command can be used to determine is the Facility MEP is collecting stats. The no form of the command disables and deletes the counters for this SAP, Binding or facility.
Default	no collect-imm-stats

csf-enable

Syntax	[no] csf-enable
Context	config>lag>eth-cfm-mep config>port>ethernet>eth-cfm>mep
Description	This command configures the reception of Client Signal Fail (CSF) message parameters.

multiplier

Syntax	multiplier <i>multiplier-value</i> no multiplier
Context	config>port>ethernet>eth-cfm>mep>csf-enable config>lag>eth-cfm-mep>csf-enable
Description	This command configures the multiplier used for timing out the CSF.
Parameters	<i>multiplier-value</i> — Specifies the multiplier used for timing out CSF.
Values	0.0, 2.0 to 30.0
Default	3.5

eth-bn

Syntax	eth-bn
Context	config>port>ethernet>eth-cfm>mep
Description	This command enables the context to configure Ethernet Bandwidth Notification (ETH-BN) message handling.

receive

Syntax	[no] receive
Context	config>port>ethernet>eth-cfm>mep>eth-bn
Description	This command enables the reception and processing of eth-bn messages and the retrieval and processing of the current bandwidth field for inclusion in dynamic egress rate adjustments. The received rate is an Layer 2 rate, and is expected to be in Mb/s. If this rate is a link rate (including preamble, start frame delimiter, and inter-frame gap), this would require the use of network egress queue groups (configured in the configure qos queue-group-templates egress queue-group "qg1" queue 1 packet-byte-offset add 20). The packet-byte-offset is not supported for default network queues.
Default	no receive

rx-update-pacing

Syntax	rx-update-pacing <i>seconds</i>
Context	config>port>ethernet>eth-cfm>mep>eth-bn

Description	This command sets the pace for update messages to and from the eth-cfm subsystem to the QoS subsystem. The most recent update messages are held by the ETH-CFM subsystem, but the most recent update is held until the expiration of the pacing timer.
Parameters	<i>seconds</i> — The time to wait before sending subsequent updates (in seconds).
Values	1 to 600
Default	5

eth-test-enable

Syntax	[no] eth-test-enable
Context	config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep config>router>if>eth-cfm>mep
Description	For this test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands: oam eth-cfm eth-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [priority <i>priority</i>] [data-length <i>data-length</i>] The no form of the command disables eth-test capabilities.

bit-error-threshold

Syntax	bit-error-threshold <i>bit-errors</i>
Context	config>eth-ring>path>eth-cfm>mep config>eth-tunnel>path>eth-cfm>mep>eth-test-enable config>lag>eth-cfm>mep
Description	This command specifies the lowest priority defect that is allowed to generate a fault alarm.
Default	bit-error-threshold 1
Parameters	<i>bit-errors</i> — Specifies the lowest priority defect.
Values	0 to 11840

test-pattern

Syntax	test-pattern {all-zeros all-ones} [crc-enable] no test-pattern
Context	config>lag>eth-cfm>mep>eth-test

	<pre>config>port>ethernet>eth-cfm>mep>eth-test config>router>if>eth-cfm>mep>eth-test</pre>
Description	This command specifies the test pattern of the ETH-TEST frames. This does not have to be configured the same on the sender and the receiver. The no form of the command reverts to the default values.
Default	test-pattern all-zeros
Parameters	all-zeros — Specifies to use all zeros in the test pattern. all-ones — Specifies to use all ones in the test pattern. crc-enable — Generates a CRC checksum.

grace

Syntax	grace
Context	<pre>config>eth-tunnel>path>eth-cfm>mep config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep</pre>
Description	This command enables the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters.

eth-ed

Syntax	eth-ed
Context	<pre>config>eth-tunnel>path>eth-cfm>mep>grace config>lag>eth-cfm>mep>grace config>port>ethernet>eth-cfm>mep>grace</pre>
Description	This command enables the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters.

max-rx-defect-window

Syntax	max-rx-defect-window <i>seconds</i> no max-rx-defect-window
Context	<pre>config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed config>lag>eth-cfm>mep>grace>eth-ed config>port>ethernet>eth-cfm>mep>grace>eth-ed</pre>

Description	This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter. The no form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used.
Default	no max-rx-defect-window
Parameters	<i>seconds</i> — Specifies the duration, in seconds, of the maximum expected defect window. Values 1 to 86400

priority

Syntax	priority <i>priority</i> no priority
Context	config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed config>lag>eth-cfm>mep>grace>eth-ed config>port>ethernet>eth-cfm>mep>grace>eth-ed
Description	This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC. The no form of the command disables the local priority configuration and sets the priority to the ccm-ltm-priority associated with this MEP.
Default	no priority
Parameters	<i>priority</i> — Specifies the priority bit. Values 0 to 7

rx-eth-ed

Syntax	[no] rx-eth-ed
Context	config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed config>lag>eth-cfm>mep>grace>eth-ed config>port>ethernet>eth-cfm>mep>grace>eth-ed
Description	This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP. The no form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP.
Default	rx-eth-ed

tx-eth-ed

Syntax	[no] tx-eth-ed
Context	config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed config>lag>eth-cfm>mep>grace>eth-ed config>port>ethernet>eth-cfm>mep>grace>eth-ed
Description	This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards. The config>eth-cfm>system>grace-tx-enable command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted. The no form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP.
Default	no tx-eth-ed

eth-vsm-grace

Syntax	eth-vsm-grace
Context	config>eth-tunnel>path>eth-cfm>mep>grace config>lag>eth-cfm>mep>grace config>port>ethernet>eth-cfm>mep>grace
Description	This command enables the context to configure Nokia ETH-CFM Grace functional parameters.

rx-eth-vsm-grace

Syntax	[no] rx-eth-vsm-grace
Context	config>eth-tunnel>path>eth-cfm>mep>grace>eth-vsm-grace config>lag>eth-cfm>mep>grace>eth-vsm-grace config>port>ethernet>eth-cfm>mep>grace>eth-vsm-grace
Description	This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP. The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect. The no form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.

Default rx-eth-vsm-grace

tx-eth-vsm-grace

Syntax [no] tx-eth-vsm-grace

Context config>eth-tunnel>path>eth-cfm>mep>grace>eth-vsm-grace
config>lag>eth-cfm>mep>grace>eth-vsm-grace
config>port>ethernet>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **config>eth-cfm>system>grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.

Default tx-eth-vsm-grace

low-priority-defect

Syntax low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}

Context config>lag>eth-cfm>mep>eth-test
config>port>ethernet>eth-cfm>mep>eth-test

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm. This setting is also used to determine the fault state of the MEP which, when enabled to do so, causes a network reaction.

Default low-priority-defect macRemErrXcon

Parameters allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon — Specifies the lowest priority defect.

Values allDefDefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
macRemErrXcon
Only DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM

remErrXconOnly DefRemoteCCM, DefErrorCCM, and
DefXconCCM
errXconOnly DefErrorCCM and DefXconCCM
xconOnly DefXconCCM; or
noXconNo defects DefXcon or lower are to be reported

mac-address

Syntax	mac-address <i>mac-address</i> no mac-address
Context	config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep config>router>if>eth-cfm>mep
Description	This command specifies the MAC address of the MEP. The no form of the command reverts to the MAC address of the MEP back to the default, that of the port, since this is SAP based.
Default	no mac-address
Parameters	<i>mac-address</i> — Specifies the MAC address of the MEP. Values 6-byte unicast mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) of the MEP. Using the all zeros address is equivalent to the no form of this command.

one-way-delay-threshold

Syntax	one-way-delay-threshold <i>seconds</i>
Context	config>eth-tunnel>path>eth-cfm>mep config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep
Description	This command enables one way delay threshold time limit.
Default	one-way-delay-threshold 3
Parameters	<i>seconds</i> — Specifies the value, in seconds, for the threshold. Values 0 to 600

facility-fault

Syntax	[no] facility-fault

Context	config>lag>eth-cfm>mep config>port>ethernet>eth-cfm>mep
Description	Allows the facility MEP to move from alarming only to network actionable function. This means a facility MEP will not merely report the defect conditions but will be able to action based on the transition of the MEP state. Without this command the facility MEP will only monitor and report and conditions of the MEP do not affect related services.
Default	no facility-fault

tunnel-fault

Syntax	tunnel-fault {accept ignore}
Context	config>service>vpls>eth-cfm config>service>vpls>sap>eth-cfm config>service>epipe>eth-cfm config>service>epipe>sap>eth-cfm config>service>ipipe>eth-cfm config>service>ipipe>sap>eth-cfm config>service>ies>eth-cfm config>service>ies>if>sap>eth-cfm config>service>ies>sub-if>grp-if>sap>eth-cfm config>service>vprn>eth-cfm config>service>vprn>if>sap>eth-cfm config>service>vprn>sub-if>grp-if>sap>eth-cfm
Description	Allows the individual service SAPs to react to changes in the tunnel MEP state. When tunnel-fault accept is configured at the service level, the SAP will react according to the service type, Epipe will set the operational flag and VPLS, IES and VPRN SAP operational state will become down on failure or up on clear. This command triggers the OAM mapping functions to mate SAPs and bindings in an Epipe service as well as setting the operational flag. If AIS generation is the requirement for the Epipe services this command is not required. See the ais-enable command under the config>service>epipe>sap>eth-cfm>ais-enable context for more details. This works in conjunction with the tunnel-fault accept on the individual SAPs. Both must be set to accept to react to the tunnel MEP state. By default the service level command is "ignore" and the SAP level command is "accept". This means simply changing the service level command to "accept" will enable the feature for all SAPs. This is not required for Epipe services that only wish to generate AIS on failure.
Default	tunnel-fault ignore (Service Level) tunnel-fault accept (SAP Level for Epipe and VPLS)
Parameters	accept — Shares fate with the facility tunnel MEP. ignore — Does not share fate with the facility tunnel MEP.

2.20.2.29 Multi-Chassis Redundancy Commands

redundancy

Syntax **redundancy**

Context config

Description This command allows the user to perform redundancy operations.

Associated commands include the following in the **admin>redundancy** context:

- **force-switchover** - Forces a switchover to the standby CPM/CFM card.
- **now** - Switch to standby CPM/CFM.

Switching to the standby displays the following message.

WARNING: Configuration and/or Boot options may have changed since the last save.

Are you sure you want to switchover (y/n)?

- **synchronize** - Synchronizes the secondary CPM/CFM.

synchronize

Syntax **synchronize {boot-env | config}**

Context config>redundancy

Description This command performs a synchronization of the standby CPM/CFM's images and/or config files to the active CPM/CFM. Either the **boot-env** or **config** parameter must be specified.

In the **config>redundancy** context, this command performs an automatically triggered standby CPM/CFM synchronization.

When the standby CPM/CFM takes over operation following a failure or reset of the active CPM/CFM, it is important to ensure that the active and standby CPM/CFMs have identical operational parameters. This includes the saved configuration, CPM and IOM images. This includes the saved configuration, CPM and IOM images. This includes the saved configuration and CFM images.

The active CPM/CFM ensures that the active configuration is maintained on the standby CPM/CFM. However, to ensure smooth operation under all circumstances, runtime images and system initialization configurations must also be automatically synchronized between the active and standby CPM/CFM.

If synchronization fails, alarms and log messages that indicate the type of error that caused the failure of the synchronization operation are generated. When the error condition ceases to exist, the alarm is cleared.

Only files stored on the router are synchronized. If a configuration file or image is stored in a location other than on a local compact flash, the file is not synchronized (for example, storing a configuration file on an FTP server).

Default	config
Parameters	boot-env — Synchronizes all files required for the boot process (loader, BOF, images, and configuration files. config — Synchronizes only the primary, secondary, and tertiary configuration files.

bgp-multi-homing

Syntax	bgp-multi-homing
Context	config>redundancy
Description	This command configures BGP multi-homing parameters.

boot-timer

Syntax	boot-timer <i>seconds</i> no boot-timer
Context	config>redundancy>bgp-mh
Description	This command specifies how long the service manager waits after a node reboot before running the MH procedures. The boot-timer value should be configured to allow for the BGP sessions to come up and for the NLRI information to be refreshed/exchanged. The boot-timer is activated after the no shutdown command for a MH site executed from configuration. Upon activation, the boot-timer is compared with the system up-time for the node. If the boot timer is higher than the up-time, then the service manager waits for the boot-timer-sys-up-time, then starts the site-activation-timer. The no form of this command sets the value to 10.
Default	no boot-timer
Parameters	<i>seconds</i> — Specifies the timer, in seconds. Values 1 to 600

site-activation-timer

Syntax	site-activation-timer <i>seconds</i> no site-activation-timer

Context	config>redundancy>bgp-mh
Description	This command defines the amount of time the service manager will keep the local sites in standby status, waiting for BGP updates from remote PEs before running the DF election algorithm to decide whether the site should be unblocked. The timer is started when one of the following event occurs only if the site is operationally up: • Manual site activation using “no shutdown” at site-id level or at member object(s) level (for example, SAP(s) or PW(s)) • Site activation after a failure The no form of this command sets the value to 2.
Default	no site-activation-timer
Parameters	<i>seconds</i> — Specifies the timer, in seconds. Values 1 to 100

site-min-down-timer

Syntax	site-min-down-timer <i>seconds</i> no site-min-down-timer
Context	config>redundancy>bgp-multi-homing
Description	This command configures the BGP multi-homing site minimum down time. When set to a non-zero value, if the site goes operationally down it will remain operationally down for at least the length of time configured for the site-min-down-timer, regardless of whether other state changes would have caused it to go operationally up. This timer is restarted every time that the site transitions from up to down. The above operation is optimized in the following circumstances: • If the site goes down on the designated forwarder but there are no BGP multi-homing peers with the same site in an UP state, then the site-min-down-timer is not started and is not used. • If the site goes down on the designated forwarder but there are no active BGP multi-homing peers, then the site-min-down-timer is not started and is not used. • If the site-min-down-timer is active and a BGP multi-homing update is received from the designated forwarder indicating its site has gone down, the site-min-down-timer is immediately terminated and this PE becomes the designated forwarder if the BGP multi-homing algorithm determines it should be the designated forwarder. The no form of the command reverts to a value of 0.
Default	no site-min-down-timer

Parameters	<i>seconds</i> — Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.
Values	1 to 100

multi-chassis

Syntax	multi-chassis
Context	config>redundancy
Description	This command enables the context to configure multi-chassis parameters.

peer

Syntax	peer <i>ip-address</i> [create] no peer <i>ip-address</i>										
Context	config>redundancy>multi-chassis										
Description	Use this command to configure up to 20 multi-chassis redundancy peers. Note that it is only for mc-lag (20) not for mc-sync (4).										
Parameters	<i>ip-address</i> — Specifies the IP address.										
Values	<table><tr><td>ipv4-address:</td><td>a.b.c.d</td></tr><tr><td>ipv6-address:</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr><tr><td></td><td>x:x:x:x:x:d.d.d.d</td></tr><tr><td></td><td>x:-[0 to FFFF]H</td></tr><tr><td></td><td>d: [0 to 255]D</td></tr></table>	ipv4-address:	a.b.c.d	ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)		x:x:x:x:x:d.d.d.d		x:-[0 to FFFF]H		d: [0 to 255]D
ipv4-address:	a.b.c.d										
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)										
	x:x:x:x:x:d.d.d.d										
	x:-[0 to FFFF]H										
	d: [0 to 255]D										
	create — Specifies to create the peer.										

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>redundancy>multi-chassis>peer
Description	This command configures the authentication key used between this node and the multi-chassis peer. The authentication key can be any combination of letters or numbers.
Default	no authentication-key

- Parameters**
- authentication-key* — Specifies the authentication key. Allowed values are any string up to 20 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
 - hash-key* — The hash key. The key can be any combination of ASCII characters up to 33 (hash1-key) or 65 (hash2-key) characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").
 - hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified
 - hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

2.20.2.29.1 MC Endpoint Commands

mc-endpoint

- Syntax** [no] mc-endpoint
- Context** config>redundancy>multi-chassis>peer
- Description** This command specifies that the endpoint is multi-chassis. This value should be the same on both MC-EP peers for the pseudowires that must be part of the same group.
- The **no** form of this command removes the endpoint from the MC-EP. Single chassis behavior applies.
- Default** no mc-endpoint

bfd-enable

- Syntax** [no] bfd-enable
- Context** config>redundancy>multi-chassis>peer>mc-ep
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
config>router>rsvp

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command disables BFD.

Default no bfd-enable

boot-timer

Syntax **boot-timer** *interval*
no boot-timer

Context config>redundancy>multi-chassis>peer>mc-ep

Description This command configures the boot timer interval. This command applies only when the node reboots. It specifies the time the MC-EP protocol keeps trying to establish a connection before assuming a failure of the remote peer. This is different from the keep-alives mechanism which is used just after the peer-peer communication was established. After this time interval passed all the mc-endpoints configured under services will revert to single chassis behavior, activating the best local PW.

The **no** form of this command sets the interval to default.

Default no boot-timer

Parameters *interval* — Specifies the boot timer interval.

Values 1 to 600

hold-on-neighbor-failure

Syntax **hold-on-neighbor-failure** *multiplier*
no hold-on-neighbor-failure

Context config>redundancy>multi-chassis>peer>mc-ep

Description This command specifies the number of keep-alive intervals that the local node will wait for packets from the MC-EP peer before assuming failure. After this time interval passed the all the mc-endpoints configured under services will revert to single chassis behavior, activating the best local pseudowire.

The **no** form of this command sets the multiplier to default value

Default no hold-on-neighbor-failure

Parameters *multiplier* — Specifies the hold time applied on neighbor failure.
Values 2 to 25

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

Context config>redundancy>multi-chassis>peer>mc-ep

Description This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-EP when bfd is not enabled or is down. These fast keep-alive messages are used to determine remote-node failure and the interval is set in deci-seconds.

The **no** form of this command sets the interval to default value

Default no keep-alive-interval

Parameters *interval* — The time interval expressed in tenths of a second.
Values 5 to 500

passive-mode

Syntax **[no] passive-mode**

Context config>redundancy>multi-chassis>peer>mc-ep

Description This command configures the passive mode behavior for the MC-EP protocol. When in passive mode the MC-EP pair will be dormant until two of the pseudowires in a MC-EP will be signaled as active by the remote PEs, being assumed that the remote pair is configured with regular MC-EP. As soon as more than one pseudowire is active, dormant MC-EP pair will activate. It will use the regular exchange to select the best pseudowire between the active ones and it will block the Rx and Tx directions of the other pseudowires.

The **no** form of this command will disable the passive mode behavior.

Default no passive-mode

system-priority

Syntax **system-priority** *value*
no system-priority

Context config>redundancy>multi-chassis>peer>mc-ep

Description	This command allows the operator to set the system priority. The peer configured with the lowest value is chosen to be the Master. If more than one peer has the same lowest system-priority value, then the one with the lowest system-id (chassis MAC address) is chosen as the Master. The no form of this command sets the system priority to 0.
Default	no system-priority
Parameters	<i>value</i> — Specifies the priority assigned to the local MC-EP peer. Values 1 to 255

2.20.2.29.2 MC LAG Commands

mc-lag

Syntax	[no] mc-lag
Context	config>redundancy>multi-chassis>peer
Description	This command enables the context to configure multi-chassis LAG operations and related parameters. The no form of this command administratively disables multi-chassis LAG. MC-LAG can be issued only when mc-lag is shutdown.
Default	no mc-lag

hold-on-neighbor-failure

Syntax	hold-on-neighbor-failure <i>multiplier</i> no hold-on-neighbor-failure
Context	config>redundancy>multi-chassis>peer>mc-lag
Description	This command specifies the interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure. This delay in switch-over operation is required to accommodate different factors influencing node failure detection rate, such as IGP convergence, or HA switch-over times and to prevent the standby node to take action prematurely. The no form of this command sets this parameter to default value.
Default	hold-on-neighbor-failure 3

Parameters *multiplier* — The time interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure.

Values 2 to 25

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-LAG. These keep-alive messages are used to determine remote-node failure and the interval is set in deci-seconds.

The **no** form of this command sets the interval to default value.

Default keep-alive-interval 10

Parameters *interval* — The time interval expressed in tenths of a second.

Values 5 to 500

lag

Syntax **lag** *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *remote-lag-id*] **system-priority** *system-priority* **source-bmac-lsb** **use-lacp-key**
lag *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *remote-lag-id*] **system-priority** *system-priority* **source-bmac-lsb** *MAC-Lsb*
lag *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *remote-lag-id*] **system-priority** *system-priority*
lag *lag-id* [**remote-lag** *remote-lag-id*]
no lag *lag-id*

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command defines a LAG which is forming a redundant-pair for MC-LAG with a LAG configured on the given peer. The same LAG group can be defined only in the scope of 1 peer. In order MC-LAG to become operational, all parameters (**lacp-key**, **system-id**, **system-priority**) must be configured the same on both nodes of the same redundant pair.

The partner system (the system connected to all links forming MC-LAG) will consider all ports using the same **lacp-key**, **system-id**, **system-priority** as the part of the same LAG. In order to achieve this in MC operation, both redundant-pair nodes have to be configured with the same values. In case of the mismatch, MC-LAG is kept in oper-down status.

Note that the correct CLI command to enable MC LAG for a LAG in **standby-signaling power-off mode** is **lag lag-id [remote-lag remote-lag-id]**. In the CLI help output, the first three forms are used to enable MC LAG for a LAG in LACP mode. MC LAG is disabled (regardless of the mode) for a given LAG with **no lag lag-id**.

- Parameters**
- lag-id* — The LAG identifier, expressed as an integer. Specifying the *lag-id* allows the mismatch between lag-id on redundant-pair. If no **lag-id** is specified it is assumed that neighbor system uses the same *lag-id* as a part of the given MC-LAG. If no matching MC-LAG group can be found between neighbor systems, the individual LAGs will operate as usual (no MC-LAG operation is established).
- Values** 1 to 800
- admin-key* — Specifies a 16 bit key that needs to be configured in the same manner on both sides of the MC-LAG in order for the MC-LAG to come up.
- Values** 1 to 65535
- system-id* — Specifies a 6 byte value expressed in the same notation as MAC address.
- Values** xx:xx:xx:xx:xx:xx - xx [00 to FF]
- remote-lag-id* — Specifies the LAG ID on the remote system.
- Values** 1 to 800
- system-priority* — Specifies the system priority to be used in the context of the MC-LAG. The partner system will consider all ports using the same **lacp-key**, **system-id**, and **system-priority** as part of the same LAG.
- Values** 1 to 65535
- MAC-Lsb* — Configures the last 16 bit of the MAC address to be used for all traffic ingressing the MC-LAG link(s) or if use-lacp-key option is used, it will only copy the value of lacp-key (redundancy multi-chassis mc-lag lag lacp-key admin-key). The command will fail if the *value* is the same with any of the following configured attributes:
- Source-bmac-lsb assigned to other MC-LAG ports.
 - lsb 16 bits value for the source-bmac configured at chassis or BVPLS level
- The first 32 bits will be copied from the source BMAC of the BVPLS associated with the IVPLS for a specific IVPLS SAP mapped to the MC-LAG. The BVPLS source BMAC can be provisioned for each BVPLS or can be inherited from the chassis PBB configuration.
- Values** 1 to 65535 or xx-xx or xx:xx

peer-name

Syntax **peer-name** *name*
no peer-name

Context config>redundancy>multi-chassis>peer

Description	This command specifies a peer name.
Default	no peer-name
Parameters	<i>name</i> — Specifies the string up to 32 characters. Any printable, seven-bit ASCII characters can be used within the string. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

source-address

Syntax	source-address <i>ip-address</i> no source-address
Context	config>redundancy>multi-chassis>peer
Description	This command specifies the source address used to communicate with the multi-chassis peer.
Parameters	<i>ip-address</i> — Specifies the source address used to communicate with the multi-chassis peer.
Values	ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D

sync

Syntax	[no] sync
Context	config>redundancy>multi-chassis>peer
Description	This command enables the context to configure synchronization parameters.
Default	no sync

igmp

Syntax	[no] igmp
Context	config>redundancy>multi-chassis>peer>sync
Description	This command specifies whether IGMP protocol information should be synchronized with the multi-chassis peer.

Default no igmp

igmp-snooping

Syntax [no] igmp-snooping

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether IGMP snooping information should be synchronized with the multi-chassis peer.

Default no igmp-snooping

mld

Syntax [no] mld

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether MLD protocol information should be synchronized with the multi-chassis peer.

Default no mld

mld-snooping

Syntax [no] mld-snooping

Context config>redundancy>multi-chassis>peer>sync

Description This command is not supported. It is not blocked for backwards-compatibility reasons but has no effect on the system if configured.

Default no mld-snooping

pim-snooping

Syntax pim-snooping [saps] [spoke-sdps]
no pim-snooping

Context config>redundancy>multi-chassis>peer>sync

Description This command specifies whether PIM snooping for IPv4 information should be synchronized with the multi-chassis peer. Entering only **pim-snooping** (without any parameter) results in the synchronization being applicable only to SAPs.

Default	no pim-snooping
Parameters	saps — Specifies that SAPs are to be synchronized with the multi-chassis peer relating to <i>sync-tags</i> configured on ports. This is the default when no parameters are included. spoke-sdp — Specifies that spoke SDPs are to be synchronized with the multi-chassis peer relating to <i>sync-tags</i> configured on ports.

port

Syntax	port <i>port-id</i> [sync-tag <i>sync-tag</i>] [create] no port <i>port-id</i>
Context	config>redundancy>multi-chassis>peer>sync
Description	This command specifies the port to be synchronized with the multi-chassis peer and a synchronization tag to be used while synchronizing this port with the multi-chassis peer.
Parameters	<i>port-id</i> — Specifies the port to be synchronized with the multi-chassis peer.

Values

slot/mda/port		
lag-id	lag-id	
	lag	keyword
	id	1 to 800
pw-id	pw-id	
	pw	keyword
	id	1 to 10239
eth-sat-id (7950 XRS only)	esat-id/slot/port	
	esat	keyword
	id	1 to 20

sync-tag — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer. The tag can be up to 32 characters long.

create — Mandatory while creating an entry.

range

Syntax	range <i>encap-range</i> sync-tag <i>sync-tag</i> no range <i>encap-range</i>
Context	config>redundancy>multi-chassis>peer>sync>port

- Description** This command configures a range of encapsulation values.
- Parameters** *encap-range* — Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer.

Values

Dot1Q	start-tag-end-tag
start-tag	0 to 4094
end-tag	0 to 4094
QinQ	qtag1.start-qtag2-qtag1.end-qtag2-start-qtag1.*-end-qtag1.*
qtag1	1 to 4094
start-qtag1	1 to 4094
en-qtag1	1 to 4094
start-qtag2	0 to 4094
end-qtag2	0 to 4094

sync-tag — Specifies a synchronization tag up to 32 characters to be used while synchronizing this encapsulation value range with the multi-chassis peer.

srrp

- Syntax** [no] srrp
- Context** config>redundancy>multi-chassis>peer>sync
- Description** This command specifies whether subscriber routed redundancy protocol (SRRP) information should be synchronized with the multi-chassis peer.
- Default** no srrp

sub-mgmt

- Syntax** sub-mgmt [ipoe] [pppoe]
no sub-mgmt
- Context** config>redundancy>multi-chassis>peer>sync
- Description** This command specifies whether subscriber management information should be synchronized with the multi-chassis peer.
- Default** no sub-mgmt
- Parameters** **ipoe** — Specifies to synchronize IPoE subscribers. The use of the keyword must match on both nodes, otherwise the subscriber synchronization fails.

pppoe — Specifies to synchronize PPPoE subscribers. The use of the keyword must match on both nodes, otherwise the subscriber synchronization fails.

2.20.2.29.3 Multi-Chassis Ring Commands

mc-ring

Syntax	[no] mc-ring
Context	config>redundancy>mc>peer config>redundancy>multi-chassis>peer>sync
Description	This command enables the context to configure the multi-chassis ring parameters.
Default	no mc-ring

ring

Syntax	ring sync-tag [create] no ring sync-tag
Context	config>redundancy>mc>peer>mcr
Description	This command configures a multi-chassis ring.
Parameters	<i>sync-tag</i> — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer. The tag can be up to 32 characters. create — Keyword used to create the multi-chassis peer ring instance. The create keyword requirement can be enabled/disabled in the environment>create context.

in-band-control-path

Syntax	in-band-control-path
Context	config>redundancy>mc>peer>mcr>ring
Description	This command enables the context to configure multi-chassis ring inband control path parameters.

dst-ip

Syntax	dst-ip ip-address

no dst-ip

Context	config>redundancy>mc>peer>mcr>ring>in-band-control-path config>redundancy>mc>peer>mcr>node>cv
Description	This command specifies the destination IP address used in the inband control connection. If the address is not configured, the ring cannot become operational.
Default	no dst-ip
Parameters	<i>ip-address</i> — Specifies the destination IP address.
Values	a.b.c.d (no multicast address)

interface

Syntax	interface <i>ip-int-name</i> no interface
Context	config>redundancy>mc>peer>mcr>ring>in-band-control-path
Description	This command specifies the name of the IP interface used for the inband control connection. If the name is not configured, the ring cannot become operational.
Default	no interface
Parameters	<i>ip-int-name</i> — Specifies the name of the IP interface. The name can be up to 32 characters and must start with a letter.

service-id

Syntax	service-id <i>service-id</i> no service-id
Context	config>redundancy>mc>peer>mcr>ring>ibc
Description	This command specifies the service ID if the interface used for the inband control connection belongs to a VPRN service. If not specified, the service-id is zero and the interface must belong to the Base router. This command supersedes the configuration of a service name. The no form of the command removes the service id from the IBC configuration.
Default	no service-id
Parameters	<i>service-id</i> — Specifies a service ID or an existing service name.
Values	1 to 2147483647 - Only supported in 'classic' configuration-mode (configure system management-interface configuration-mode classic)

service-name - A string up to 64 characters.

service-name

Syntax	service-name <i>service-name</i> no service-name
Context	config>redundancy>mc>peer>mcr>ring>ibc
Description	This command specifies the service name if the interface used for the inband control connection belongs to a VPRN service. If not specified the interface must belong to the Base router. This command supersedes the configuration of a service ID. The no form of the command removes the service name from the IBC configuration.
Default	no service-name
Parameters	<i>service-name</i> — Specifies a service name. Values A string up to 64 characters.

path-b

Syntax	[no] path-b
Context	config>redundancy>mc>peer>mcr>ring
Description	This command specifies the set of upper-VLAN IDs associated with the SAPs that belong to path B with respect to load-sharing. All other SAPs belong to path A.
Default	If not specified, the default is an empty set.

range

Syntax	[no] range <i>vlan-range</i>
Context	config>redundancy>mc>peer>mcr>ring>path-b config>redundancy>mc>peer>mcr>ring>path-excl
Description	This command configures a MCR b-path VLAN range.
Parameters	<i>vlan-range</i> — Specifies the VLAN range. Values [0 to 4094] to [0 to 4094] [0 to 4094] to * * to *

path-excl

Syntax	[no] path-excl
Context	config>redundancy>mc>peer>mcr>ring
Description	This command specifies the set of upper-VLAN IDs associated with the SAPs that are to be excluded from control by the multi-chassis ring.
Default	If not specified, the default is an empty set.

ring-node

Syntax	ring-node <i>ring-node-name</i> [create] no ring-node <i>ring-node-name</i>
Context	config>redundancy>mc>peer>mcr>ring
Description	This command specifies the unique name of a multi-chassis ring access node.
Parameters	<i>ring-node-name</i> — Specifies the unique name of a multi-chassis ring access node. The name can be up to 32 characters. create — Keyword used to create the ring node instance. The create keyword requirement can be enabled or disabled in the environment>create context.

connectivity-verify

Syntax	connectivity-verify
Context	config>redundancy>mc>peer>mcr>ring>ring-node
Description	This command enables the context to configure node connectivity check parameters.

interval

Syntax	interval <i>interval</i> no interval
Context	config>redundancy>mc>peer>mcr>node>cv
Description	This command specifies the polling interval of the ring-node connectivity verification of this ring node.
Default	interval 5

Parameters *interval* — Specifies the polling interval, in minutes.
Values 1 to 6000

service-id

Syntax **service-id** *service-id*
no service-id

Context config>redundancy>mc>peer>mcr>node>cv

Description This command specifies the service ID of the SAP used for the ring-node connectivity verification of this ring node. This command supersedes the configuration of a service name.

The **no** form of the command removes the service id from the CV configuration.

Default no service-id

Parameters *service-id* — Specifies a service ID or an existing service name.
Values 1 to 2147483647 - Only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**)
service-name - A string up to 64 characters.

service-name

Syntax **service-name** *service-name*
no service-name

Context config>redundancy>mc>peer>mcr>node>cv

Description This command specifies the service name of the SAP used for ring-node connectivity verification of this ring node. This command supersedes the configuration of a service ID.

The **no** form of the command removes the service name from the CV configuration.

Default no service-id

Parameters *service-name* — Specifies a service name.
Values A string up to 64 characters.

src-ip

Syntax **src-ip** *ip-address*
no src-ip

Context	config>redundancy>mc>peer>mcr>node>cv
Description	This command specifies the source IP address used in the ring-node connectivity verification of this ring node.
Default	no src-ip
Parameters	<i>ip-address</i> — Specifies the source IP address. Values a.b.c.d (no multicast address)

src-mac

Syntax	src-mac <i>ieee-address</i> no src-mac
Context	config>redundancy>mc>peer>mcr>node>cv
Description	This command specifies the source MAC address used for the Ring-Node Connectivity Verification of this ring node. A value of all zeros (000000000000 H (0:0:0:0:0:0)) specifies that the MAC address of the system management processor (CPM) is used.
Default	no src-mac
Parameters	<i>ieee-address</i> — Specifies the source MAC address. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

vlan

Syntax	vlan <i>vlan-encap</i> no vlan									
Context	config>redundancy>mc>peer>mcr>node>cv									
Description	This command specifies the VLAN tag used for the Ring-node Connectivity Verification of this ring node. It is only meaningful if the value of service ID is not zero. A zero value means that no VLAN tag is configured.									
Default	no vlan									
Parameters	<i>vlan-encap</i> — Specifies the VLAN tag. Values <table><tr><td>vlan-encap:</td><td>dot1q</td><td>qtag, *</td></tr><tr><td></td><td>qinq</td><td>qtag1.qtag2, qtag1.*, 0.*</td></tr><tr><td></td><td>qtag</td><td>0 to 4094</td></tr></table>	vlan-encap:	dot1q	qtag, *		qinq	qtag1.qtag2, qtag1.*, 0.*		qtag	0 to 4094
vlan-encap:	dot1q	qtag, *								
	qinq	qtag1.qtag2, qtag1.*, 0.*								
	qtag	0 to 4094								

qtag1 1 to 4094
qtag2 0 to 4094

2.20.2.30 Forwarding Plane Tools Commands

cpm

- Syntax** **cpm**
- Context** tools>dump>mcast-path-mgr
- Description** This command dumps multicast path manager CPM information.
- Output** The following output is an example of CPM information.

Sample Output

```
*A:PE-1# tools dump mcast-path-mgr cpm
=====
Slot: 5 Complex: 0
=====
PATH:
Type SGs      InUseBW  AvailBW  TotalBW  ID  SGs      InUseBW  AvailBW  TotalBW
P      7      28040    -        -    17    7      28040    1971960  2000000
P      7      28047    -        -    16    7      28047    1971953  2000000
P      7      28047    -        -    18    7      28047    1971953  2000000
P      6      24075    -        -    19    6      24075    1975925  2000000
P      6      24075    -        -    20    6      24075    1975925  2000000
P      6      24097    -        -    21    6      24097    1975903  2000000
P      7      28076    -        -    22    7      28076    1971924  2000000
P      7      28043    -        -    23    7      28043    1971957  2000000
P      7      28018    -        -    24    7      28018    1971982  2000000
P      7      28036    -        -    25    7      28036    1971964  2000000
P      7      28036    -        -    26    7      28036    1971964  2000000
P      6      24042    -        -    27    6      24042    1975958  2000000
P      6      24053    -        -    28    6      24053    1975947  2000000
P      7      28047    -        -    29    7      28047    1971953  2000000
P      7      28080    -        -    30    7      28080    1971920  2000000
S      0          0    -        -    31    0          0  1800000  1800000
B      0          0    -        -    -      -          -    -        -
*A:PE-1#
```


2.21 Show, Monitor, Clear, Debug, and Tools Command Reference

2.21.1 Command Hierarchies

- [Show Commands](#)
- [Monitor Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

2.21.1.1 Show Commands

```
show
— aps [aps-id] [detail]
— card slot-number cpu [sample-period seconds]
— card slot-number fp fp-number dist-cpu-protection
— card slot-number fp fp-number policy-accounting
— card slot-number memory-pools
— card state
— card [slot-number]
— card [slot-number] detail
— card slot-number fp fp-number fwd-engine drop-reason statistics
— card slot-number fp fp-number ingress queue-group mode {access | network}
— card slot-number [detail] fp fp-number ingress queue-group queue-group-name instance
   [instance-number] mode {access | network} [statistics]
— cflowd
— chassis [environment] [power-supply] [ccm] [chassis-id]
— chassis class chassis-class
— chassis [detail]
— chassis detail [class chassis-class]
   — power-management [zone]
     — requirements [detail]
     — utilization [detail]
— datapath slot[/mda[/connector[/port]]] [mac mac-chip-number] [fp fp-number] [detail]
— elmi
   — evc [port-id [vlan vlan-id]]
   — uni [port-id]
— eth-tunnel {aps | status}
— eth-tunnel tunnel-index [path path-index] [detail]
— fwd-path-ext [fpe fpe-id]
— fwd-path-ext fpe fpe-id associations
— interface-group-handler [index]
```

- **lag** [*lag-id*] [**detail**] [**statistics**]
- **lag** [*lag-id*] **description**
- **lag** [*lag-id*] **port**
- **lag** *lag-id* **associations per-link-hash interface class** {1 | 2 | 3}
- **lag** *lag-id* **associations**
- **lag** *lag-id* **bfd**
- **lag** *lag-id* [**detail**] **eth-cfm** [*tunnel tunnel-id*]
- **lag** *lag-id* **associations per-link-hash interface**
- **lag** *lag-id* **associations link-map-profile** [*link-map-profile*] **interface**
- **lag** *lag-id* **lACP-partner**
- **lag** *lag-id* **detail lACP-partner**
- **lag** *lag-id* **link-map-profile** *link-map-profile*
- **lag** *lag-id* **per-link-hash** [**class** {1 | 2 | 3}]
- **lag** *lag-id* **associations per-link-hash sAP** [**class** {1 | 2 | 3}]
- **lag** *lag-id* **associations link-map-profile** [*link-map-profile*] **sAP**
- **lag** *lag-id* **per-link-hash port** *port-id*
- **licensing** **entitlements**
- **licensing** *port-id*
- **macsec**
 - **connectivity-association** [*ca-name*] [**detail**]
 - **mka-session** [**port** *port-id*]
 - **mka-session** [**port** *port-id*] **detail**
 - **mka-session** [**port** *port-id*] **statistics**
- **mda** *slot* [*mda*] [**detail**]
 - **qos** {**ingress** | **egress**} **buffer-allocation** [**detail**]
 - **qos** {**ingress** | **egress**} **orphaned-queues**
- **megapools** *slot-number* **fp** *forwarding-plane* **wred** [**detail**] **queue-group** *queue-group-name* [*instance instance-id*]
- **megapools** *slot-number* **fp** *forwarding-plane*
- **megapools** *slot-number* **fp** *forwarding-plane* **wred** [**detail**] [**service-id** *service-id*]
- **multilink-bundle** [*bundle-id* | *bpgrp-id* | *slot/mda*] **type** {**mlppp** | **ima-grp** | **mlfr**} [**detail**]
- **multilink-bundle** {*bundle-id* | *bpgrp-id* | *slot/mda*} [**ppp** | **ima** | **mlfr**]
- **multilink-bundle** {*bundle-id* | *bpgrp-id*} **relations**
 - **ima**
 - **atm** [**detail**]
 - **connections**
 - **port-connection** [**detail**]
 - **pvc** [*vpi/vci*] [**detail**]
 - **pvp** [*vpi*] [**detail**]
 - **pvt** [*vpi.vpi*] [**detail**]
 - **ppp** [**multiclass**]
 - **relations**
- **pcm** [*pcm-slot*] [**chassis** *chassis-id*] [**detail**]
- **peq** [*peq-slot*] [**chassis** *chassis-id*] [**detail**]
- **pools** *mda-id* [*port*] [*access-app* [*pool-name* | **service** *service-id* | **queue-group** *queue-group-name* [*instance id*]]]
- **pools** *mda-id* [*port*] [*network-app* [*pool-name* | **queue-group** *queue-group-name* [*instance id*]]]
- **pools** *mda-id* [*port*] [*direction* [*pool-name* | **service** *service-id* | **queue-group** *queue-group-name* [*instance id*]]]
- **port** [*port-id*] [**statistics** [**egress-aggregate**]] [**detail**]
- **port** *port-id* **associations**
- **port** [*port-id*] **description**

- **port** *port-id* dotx1 [detail]
- **port** aps
- **port** [*port-id*] cem
- **port** *port-id* atm [detail]
- **port** *port-id* atm connections
- **port** *port-id* atm cp [*cp*] [detail]
- **port** *port-id* atm ilmi [detail]
- **port** *port-id* atm interface-connection [detail]
- **port** *port-id* atm pvc [*vpi*[/*vci*]] [detail]
- **port** *port-id* atm pvp [*vpi*] [detail]
- **port** *port-id* atm pvt [*vpi1.vpi2*] [detail]
- **port** *port-id* cisco-hldc
- **port** *port-id* ethernet [associations | detail]
- **port** *port-id* ethernet [associations | detail] efm-oam event-logs [{failure | degraded}] [{active | cleared}]
- **port** *port-id* ethernet [associations | detail] lldp [nearest-bridge | nearest-non-tpmr | nearest-customer] [remote-info] [detail]
- **port** *port-id* frame-relay [detail]
- **port** *port-id* ima-link
- **port** *port-id* macsec sub-port *sub-port-id* [statistics] [detail]
- **port** *port-id* mlfr-link
- **port** *port-id* monitor-threshold
- **port** *port-id* optical
- **port** *port-id* otu
- **port** *port-id* ppp [detail]
- **port** *port-id* queue-group [ingress | egress] [*queue-group-name*] [access | network] [*instance instance-id*]
- **port** *port-id* queue-group [ingress | egress] [*queue-group-name*] [access | network] associations [*instance instance-id*]
- **port** *port-id* queue-group [ingress | egress] *queue-group-name* [access | network] [*instance instance-id*] queue-depth [*queue queue-id*]
- **port** *port-id* queue-group [ingress | egress] [*queue-group-name*] [access | network] statistics [*instance instance-id*]
- **port** *port-id* queue-group summary
- **port** *port-id* vport *vport-name* monitor-threshold
- **port** *port-id* wavekey-table {50g | 100g}
- **port** *port-id* wavelength-table
- **port** *port-id* wavetracker
- **port-tree** *port-id*
- **redundancy**
 - **multi-chassis**
 - **all**
 - **mc-lag** peer *ip-address* [*lag lag-id*]
 - **mc-lag** [peer *ip-address* [*lag lag-id*]] statistics
 - **mc-ring** peer *ip-address* statistics
 - **mc-ring** peer *ip-address* [ring *sync-tag* [detail | statistics]]
 - **mc-ring** peer *ip-address* ring *sync-tag* ring-node [*ring-node-name* [detail | statistics]]
 - **mc-ring** global-statistics
 - **sync** [peer *ip-address*]

- **sync** peer *ip-address* detail
- **sync** [peer *ip-address*] statistics
- **wpp** peer [*ip-address*] statistics
- **sfm** [*sfm-name*] [detail]
- **sfm** [*sfm-name*] icport [down] [degraded]
- **system**
 - **ethernet**
 - **efm-oam**
 - **license** available-licenses
 - **lldp**
 - **lldp** neighbor
 - **switch-fabric** [exclude-sfm *sfm-list*]
 - **switch-fabric** high-bandwidth-multicast

2.21.1.2 Monitor Commands

For more information about monitor commands, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide* for command usage and CLI syntax.

monitor

- **card** *slot-number* **fp** *fp-number* **ingress** {**access** | **network**} **queue-group** *queue-group-name* **instance** *instance-id* [**interval** *seconds*] [**repeat** *repeat*] **policer** *policer-id* [**absolute** | **percent-rate** [*reference-rate*]]
- **port** *port-id* [*port-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] [**multiclass**]
- **port** *port-id* [*port-id*] **atm** [**aal-5** | **cp** *cp-id* | **ilmi** | **interface-connection** | **pvc** *vpi/vci* [**aal-5** | **oam**] | **pvp** *vpi* [**aal-5** | **oam**] | **pvt** *vpi1.vpi2*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
- **port** **all-ethernet-rates** [**interval** *seconds*] [**repeat** *repeat*]
 - **queue-group** *queue-group-name* **egress** **access** [**instance** *instance-id*] [**egress-queue** *egress-queue-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
 - **queue-group** *queue-group-name* **ingress** **access** **ingress-queue** *ingress-queue-id* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
 - **queue-group** *queue-group-name* **egress** **network** **instance** *instance-id* [**policer** *policer-id*] [**egress-queue** *egress-queue-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
 - **vport** *name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
 - **vport** *name* [**interval** *seconds*] [**repeat** *repeat*] **monitor-threshold**

2.21.1.3 Clear Commands

clear

- **card** *slot-number*
- **card** *slot-number* **fp** *fp-number* **hs-pool** **high-water-mark**
- **card** *slot-number* **fp** *fp-number* **hs-pool** **high-water-mark** **mid-pool** [*value*]
- **card** *slot-number* **fp** *fp-number* **hs-pool** **high-water-mark** **root-pool** [*value*]

- **card** *slot-number* **fp** *fp-number* **hs-pool** **high-water-mark** **system**
- **card** *slot-number* **fp** *fp-number* **dist-cpu-protection**
- **card** *slot-number* **soft** [**hard-reset-unsupported-mdas**]
- **card** *slot-number* **fp** *fp-number* **fwd-engine** **drop-reason** **statistics**
- **card** *slot-number* **fp** *fp-number* **ingress** **mode** {**access** | **network**} **queue-group** *group-name*
instance *instance* **statistics**
- **lag** *lag-id* **statistics**
- **lag** *lag-id* **bfd-session** **family** {**ipv4** | **ipv6**} **lag-port** *port-id*
- **mda** *mda-id* [**statistics**]
- **port** {*port-id* | *bundle-id* | *bpgrp-id* | *aps-id*} **atm**
- **port** {*port-id* | *bundle-id* | *bpgrp-id* | *aps-id*} **statistics**
- **port** *port-id* **ethernet** **dampening**
- **port** *port-id* **ethernet** **efm-oam** **events** [**local** | **remote**]
- **port** *port-id* **exp-secondary-shaper** *shaper-name* **statistics**
- **port** *port-id* **fwd-engine** **drop-reason** **statistics**
- **port** *port-id* **hs-secondary-shaper** *shaper-name* **statistics**
- **port** *port-id* **monitor-threshold**
- **port** *port-id* **hs-pool** **high-water-mark** {[**standard value**] | [**alternate value**]}
- **port** *port-id* **phy-state-change-count**
- **port** *port-id* **queue-group** *queue-group-name* [**instance** *instance-id*] **queue-depth** [**queue**
queue-id] {**ingress** | **egress**} [**access** | **network**]
- **port** *port-id* **queue-group** *queue-group-name* [**instance** *instance-id*] [**access** | **network**]
{**ingress** | **egress**} **statistics**
- **port** *port-id* **vport** [*name*] **monitor-threshold**
- **port** *port-id* **vport** *name* **statistics**

2.21.1.4 Debug Commands

- debug**
- [no] **atm**
 - **ilmi** [*port-id*]
 - **no ilmi** *port-id*
 - **cisco-hdlc** *port-id*
 - **no cisco-hdlc**
 - [no] **frame-relay**
 - [no] **frf16** *port-id*
 - **lmi** [*port-id*]
 - **no lmi**
 - **lag** [*lag-id* *lag-id*] [**port** *port-id*] [**all**]
 - **lag** [*lag-id* *lag-id*] [**port** *port-id*] [**sm**] [**pkt**] [**cfg**] [**red**] [**iom-upd**] [**port-state**] [**timers**] [**sel-logic**]
[**mc**] [**mc-pkt**]
 - **no lag** [*lag-id* *lag-id*]
 - **ppp** *port-id*
 - **no ppp** [*port-id*]

2.21.1.5 Tools Commands

tools

```

— dump
— aps aps-id [clear]
— aps mc-aps-signaling [clear]
— aps mc-aps-ppp [clear]
— eth-tunnel tunnel-index [clear]
— lag lag-id lag-id
— map-to-phy-port {ccag ccag-id | lag lag-id | eth-tunnel tunnel-index} {isid isid [end-  
isid isid] | service service-id | svc-name [end-service service-id | svc-name]}  
[summary]
— port port-id
— dwdm
— coherent
— cpr-wndw-sz-srch-info
— pcs [clear]
— rs-fec [clear]
— redundancy
— multi-chassis
— mc-ring
— srrp-sync-database [instance instance-id] [peer ip-address]
— sync-database [peer ip-address] [port port-id | lag-id] [sync-tag sync-  
tag] [application application] [detail] [type type]
— sync-database [peer ip-address] [sdp sdp-id] [sync-tag sync-tag]  
[application application] [detail] [type type]

tools
— perform
— aps
— clear aps-id {protect | working} [number number]
— clear-lockout-annexb aps-id
— exercise aps-id {protect | working}
— force aps-id {protect | working} [number number]
— lockout aps-id
— lockout-annexb aps-id
— request aps-id {protect | working}
— eth-ring
— clear ring-index
— force ring-index path {a | b}
— manual ring-index path {a | b}
— ima
— reset bundle-id
— lag
— clear-force all-mc
— clear-force lag-id lag-id [sub-group sub-group-id]
— clear-force peer-mc ip-address
— force all-mc {active | standby}
— force lag-id lag-id [sub-group sub-group-id] {active | standby}
— force peer-mc ip-address {active | standby}
— load-balance lag-id lag-id [class {1 | 2 | 3}]

```


2.21.2 Command Descriptions

- [Hardware Show Commands](#)
- [PEQ Show Commands](#)
- [APS Show Commands](#)
- [Port Show Commands](#)
- [Multilink Bundle Show Commands](#)
- [LAG Show Commands](#)
- [Monitor Commands](#)
- [Clear Commands](#)
- [Tools Commands](#)



Note: The command outputs in this chapter are examples only; actual displays may differ depending on supported functionality and user configuration.

2.21.2.1 Hardware Show Commands

chassis

Syntax	<code>chassis [environment] [power-supply]</code> <code>chassis [detail]</code> <code>chassis [environment] [chassis-id] [ccm]</code> <code>chassis detail [class chassis-class]</code> <code>chassis class chassis-class</code>
Context	show
Description	This command displays general chassis status information.
Parameters	<i>chassis-id</i> — Specifies chassis 1, 2, etc for the router chassis. <i>chassis-class</i> — Shows information related to chassis of the specified class. Values router, or eth-sat environment — Shows chassis environmental status information. Default Shows all chassis information. power-supply — Shows chassis power supply status information. Default Shows all chassis information.

detail — Displays detailed information for all chassis in the system.

ccm — Shows chassis control module information. Only supported on platforms that have CCMs.

Output See the following sections for sample outputs:

- [Output Fields: show chassis](#)
- [Sample Output: show chassis \(showing 7750 SR\)](#)
- [Sample Output: show chassis environment](#)
- [Sample Output: show chassis power-supply](#)
- [Sample Output: show chassis ccm](#)
- [Sample Output: show chassis \(showing 7950 XRS\)](#)
- [Sample Output: show chassis detail \(showing 7950 XRS\)](#)
- [Sample Output: show chassis \(showing 7750 SR with Ethernet satellites configured\)](#)

Sample Output: show chassis (showing 7750 SR)

```
A:PE-1# show chassis
=====
System Information
=====
Name                  : PE-1
Type                  : 7750 SR-12
Chassis Topology      : Standalone
Location              : (Not Specified)
Coordinates           : (Not Specified)
CLLI code             :
Number of slots       : 12
Oper number of slots  : 12
Num of faceplate ports/connectors : 97
Num of physical ports : 97
Critical LED state    : Off
Major LED state       : Off
Minor LED state       : Off
Over Temperature state : OK
Base MAC address      : 10:e8:78:b8:21:47
Admin chassis mode    : d
Oper chassis mode     : d
Fabric Speed          : 10 Gig
FP Generations        : FP3
System Profile        : none
=====
Chassis Summary
=====
Chassis  Role          Status
-----
1        Standalone    up
=====
A:PE-1#
```

Sample Output: show chassis environment

```
A:SR-12# show chassis environment
=====
Chassis 1 Detail
=====
Environment Information
  Number of fan trays : 3
  Number of fans : 6
  Fan tray number : 1
    Speed : 44 %
    Status : up
  Fan tray number : 2
    Speed : 44 %
    Status : up
  Fan tray number : 3
    Speed : 44 %
    Status : up
=====
```

Sample Output: show chassis power-supply

```
A:ALA-4# show chassis power-supply
=====
Chassis Information
=====
Power Supply Information
  Number of power supplies      : 2

  Power supply number          : 1
  Defaulted power supply type   : dc
  Status                        : up

  Power supply number          : 2
  Defaulted power supply type   : dc
  Status                        : up
=====
A:ALA-4#
```

Sample Output: show chassis ccm

```
A:7750-3# show chassis ccm
=====
Chassis Information
=====
Chassis Control Module (CCM) Information
  CCM number                    : 1
  Equipped                      : yes
  Type                          : ccm-xp

Hardware Data
  Part number                   : Sim Part#
  CLEI code                     : Sim CLEI
  Serial number                 : ccm-0
  Manufacture date              : 01012003
  Manufacturing string           : Sim MfgString ccm-0
  Manufacturing deviations      : Sim MfgDeviation ccm-0
  Administrative state          : up
  Operational state             : up
```

```

Temperature           : 32C
Temperature threshold : 75C
Time of last boot     : N/A
Current alarm state   : alarm cleared

```

```

=====
A:7750-3>

```

Sample Output: show chassis (showing 7950 XRS)

```

*A:PE-1# show chassis
=====
System Information
=====
Name           : PE-1
Type           : 7950 XRS-20
Chassis Topology : Standalone
Location       : (Not Specified)
Coordinates    : (Not Specified)
CLLI code      :
Number of slots : 24
Oper number of slots : 12
Num of faceplate ports/connectors : 314
Num of physical ports : 314
Critical LED state : Off
Major LED state  : Off
Minor LED state  : Amber
Over Temperature state : OK
Base MAC address : 24:af:4a:a3:ac:c5
FP Generations   : FP3
System Profile   : none
=====
Chassis Summary
=====
Chassis  Role          Status
-----
1        Standalone    up
=====
*A:PE-1#

```

Sample Output: show chassis detail (showing 7950 XRS)

```

*A:PE-1# show chassis detail
=====
System Information
=====
Name           : PE-1
Type           : 7950 XRS-20
Chassis Topology : Standalone
Location       : (Not Specified)
Coordinates    : (Not Specified)
CLLI code      :
Number of slots : 24
Oper number of slots : 12
Num of faceplate ports/connectors : 314
Num of physical ports : 314
Critical LED state : Off
Major LED state  : Off

```

```

Minor LED state           : Amber
Over Temperature state    : OK
Base MAC address          : 24:af:4a:a3:ac:c5
FP Generations            : FP3
System Profile            : none
=====
Chassis 1 Detail
=====
Chassis Status            : up
Chassis Role              : Standalone
Hardware Data
  Part number             :
  CLEI code               :
  Serial number           : AME4-14
  Manufacture date        :
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number :
  Time of last boot       : 2017/10/23 12:29:41
  Current alarm state     : alarm active
-----
Environment Information
Number of fan trays       : 2
Number of fans           : 16

Fan tray number          : 1
  Speed                  : 38 %
  Status                 : up
  Hardware Data
    Part number          : 3HE07119AARB01
    CLEI code            : IPUCA9U1AA
    Serial number        : NS131062838
    Manufacture date     : 03212013
    Manufacturing deviations : (Not Specified)
    Manufacturing assembly number : 82-0354-03
    Administrative state : up
    Operational state    : up
    Time of last boot    : 2017/10/23 12:27:47
    Current alarm state  : alarm cleared
    Firmware revision status : acceptable
  Hardware Resources (Power-Zone 1)
    Voltage
      Minimum            : 53.00 Volts (10/23/2017 12:30:28)
      Current            : 53.00 Volts
      Peak               : 54.00 Volts (10/23/2017 12:29:48)
    Wattage
      Minimum            : 106.00 Watts (10/23/2017 12:30:28)
      Current            : 159.00 Watts
      Peak               : 324.00 Watts (10/23/2017 12:29:48)
      Max Required       : 900.00 Watts
    Amperage
      Minimum            : 2.00 Amps (10/23/2017 12:29:56)
      Current            : 3.00 Amps
      Peak               : 6.00 Amps (10/23/2017 12:29:48)

Fan tray number          : 2
  Speed                  : 38 %
  Status                 : up
  Hardware Data
    Part number          : 3HE07119AARB01

```

```

CLEI code           : IPUCA9U1AA
Serial number       : NS152062365
Manufacture date    : 05232015
Manufacturing deviations : (Not Specified)
Manufacturing assembly number : 82-0354-03
Administrative state : up
Operational state   : up
Time of last boot   : 2017/10/23 12:27:25
Current alarm state  : alarm cleared
Firmware revision status : acceptable
Hardware Resources (Power-Zone 1)
Voltage
  Minimum           : 52.00 Volts (10/23/2017 12:32:01)
  Current            : 52.00 Volts
  Peak               : 54.00 Volts (10/23/2017 12:29:54)
Wattage
  Minimum           : 106.00 Watts (10/23/2017 12:30:05)
  Current            : 156.00 Watts
  Peak               : 371.00 Watts (10/23/2017 12:29:51)
  Max Required      : 900.00 Watts
Amperage
  Minimum           : 2.00 Amps (10/23/2017 12:29:54)
  Current            : 3.00 Amps
  Peak               : 7.00 Amps (10/23/2017 12:29:51)
-----
Power Management Information
Power Management Mode : basic
Power Safety Level    : 100%
Power Safety Alert    : 0 watts
Number of PEQs        : 12

PEQ number           : 1
PEQ Equipped Type     : apeq-dc-2000
PEQ Provisioned Type   : (Not Specified)
Power-Zone            : 1
Status                : up
Input Feed Status      : input B down
Input Power Mode       : N/A
Hardware Data
  Part number         : 3HE07114AARA01
  CLEI code           : IPUPAJHUAA
  Serial number       : NS1250G0116
  Manufacture date    : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state : up
  Operational state   : unprovisioned
  Time of last boot   : 2017/10/23 12:27:33
  Current alarm state  : alarm active
  Firmware revision status : acceptable

PEQ number           : 2
PEQ Equipped Type     : apeq-dc-2000
PEQ Provisioned Type   : (Not Specified)
Power-Zone            : 1
Status                : up
Input Feed Status      : input B down
Input Power Mode       : N/A
Hardware Data

```

```

Part number           : 3HE07114AARA01
CLEI code             : IPUPAJHUAA
Serial number         : NS1249G0022
Manufacture date      : 12202012
Manufacturing deviations : (Not Specified)
Manufacturing assembly number : 8205320107
Administrative state   : up
Operational state      : unprovisioned
Time of last boot      : 2017/10/23 12:27:26
Current alarm state    : alarm active
Firmware revision status : acceptable

PEQ number            : 3
PEQ Equipped Type      : apeq-dc-2000
PEQ Provisioned Type   : (Not Specified)
Power-Zone             : 1
Status                 : up
Input Feed Status       : input B down
Input Power Mode        : N/A
Hardware Data
  Part number          : 3HE07114AARA01
  CLEI code            : IPUPAJHUAA
  Serial number         : NS1250G0141
  Manufacture date      : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state   : up
  Operational state      : unprovisioned
  Time of last boot      : 2017/10/23 12:27:31
  Current alarm state    : alarm active
  Firmware revision status : acceptable

PEQ number            : 4
PEQ Equipped Type      : apeq-dc-2000
PEQ Provisioned Type   : (Not Specified)
Power-Zone             : 1
Status                 : up
Input Feed Status       : input B down
Input Power Mode        : N/A
Hardware Data
  Part number          : 3HE07114AARA01
  CLEI code            : IPUPAJHUAA
  Serial number         : NS1249G0201
  Manufacture date      : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state   : up
  Operational state      : unprovisioned
  Time of last boot      : 2017/10/23 12:27:26
  Current alarm state    : alarm active
  Firmware revision status : acceptable

PEQ number            : 5
PEQ Equipped Type      : apeq-dc-2000
PEQ Provisioned Type   : (Not Specified)
Power-Zone             : 1
Status                 : up
Input Feed Status       : input B down
Input Power Mode        : N/A

```

```

Hardware Data
  Part number           : 3HE07114AARA01
  CLEI code             : IPUPAJHUAA
  Serial number         : NS1250G0123
  Manufacture date      : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state   : up
  Operational state      : unprovisioned
  Time of last boot      : 2017/10/23 12:27:32
  Current alarm state    : alarm active
  Firmware revision status : acceptable

PEQ number              : 6
PEQ Equipped Type       : apeq-dc-2000
PEQ Provisioned Type     : (Not Specified)
Power-Zone              : 1
Status                  : up
Input Feed Status        : input B down
Input Power Mode         : N/A
Hardware Data
  Part number           : 3HE07114AARA01
  CLEI code             : IPUPAJHUAA
  Serial number         : NS1250G0061
  Manufacture date      : 12182012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state   : up
  Operational state      : unprovisioned
  Time of last boot      : 2017/10/23 12:27:34
  Current alarm state    : alarm active
  Firmware revision status : acceptable

PEQ number              : 7
PEQ Equipped Type       : apeq-dc-2000
PEQ Provisioned Type     : (Not Specified)
Power-Zone              : 1
Status                  : up
Input Feed Status        : input B down
Input Power Mode         : N/A
Hardware Data
  Part number           : 3HE07114AARB01
  CLEI code             : IPUPAJHUAA
  Serial number         : NS13226A310
  Manufacture date      : 06042013
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 82-0532-02
  Administrative state   : up
  Operational state      : unprovisioned
  Time of last boot      : 2017/10/23 12:27:33
  Current alarm state    : alarm active
  Firmware revision status : acceptable

PEQ number              : 8
PEQ Equipped Type       : apeq-dc-2000
PEQ Provisioned Type     : (Not Specified)
Power-Zone              : 1
Status                  : up
Input Feed Status        : input B down

```

```

Input Power Mode                : N/A
Hardware Data
  Part number                    : 3HE07114AARA01
  CLEI code                      : IPUPAJHUAA
  Serial number                  : NS1250G0152
  Manufacture date               : 12202012
  Manufacturing deviations        : (Not Specified)
  Manufacturing assembly number  : 8205320107
  Administrative state           : up
  Operational state              : unprovisioned
  Time of last boot              : 2017/10/23 12:27:29
  Current alarm state            : alarm active
  Firmware revision status       : acceptable

PEQ number                       : 9
PEQ Equipped Type                : apeq-dc-2000
PEQ Provisioned Type             : (Not Specified)
Power-Zone                      : 1
Status                          : up
Input Feed Status                : input B down
Input Power Mode                 : N/A
Hardware Data
  Part number                    : 3HE07114AARA01
  CLEI code                      : IPUPAJHUAA
  Serial number                  : NS1250G0122
  Manufacture date               : 12202012
  Manufacturing deviations        : (Not Specified)
  Manufacturing assembly number  : 8205320107
  Administrative state           : up
  Operational state              : unprovisioned
  Time of last boot              : 2017/10/23 12:27:28
  Current alarm state            : alarm active
  Firmware revision status       : acceptable

PEQ number                       : 10
PEQ Equipped Type                : apeq-dc-2000
PEQ Provisioned Type             : (Not Specified)
Power-Zone                      : 1
Status                          : up
Input Feed Status                : input B down
Input Power Mode                 : N/A
Hardware Data
  Part number                    : 3HE07114AARA01
  CLEI code                      : IPUPAJHUAA
  Serial number                  : NS1250G0146
  Manufacture date               : 12202012
  Manufacturing deviations        : (Not Specified)
  Manufacturing assembly number  : 8205320107
  Administrative state           : up
  Operational state              : unprovisioned
  Time of last boot              : 2017/10/23 12:27:25
  Current alarm state            : alarm active
  Firmware revision status       : acceptable

PEQ number                       : 11
PEQ Equipped Type                : apeq-dc-2000
PEQ Provisioned Type             : (Not Specified)
Power-Zone                      : 1
Status                          : up

```

```

Input Feed Status      : input B down
Input Power Mode       : N/A
Hardware Data
  Part number          : 3HE07114AARA01
  CLEI code            : IPUPAJHUAA
  Serial number        : NS1249G0202
  Manufacture date     : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state  : up
  Operational state    : unprovisioned
  Time of last boot    : 2017/10/23 12:27:30
  Current alarm state  : alarm active
  Firmware revision status : acceptable

PEQ number             : 12
PEQ Equipped Type      : apeq-dc-2000
PEQ Provisioned Type   : (Not Specified)
Power-Zone             : 1
Status                 : up
Input Feed Status      : input B down
Input Power Mode       : N/A
Hardware Data
  Part number          : 3HE07114AARA01
  CLEI code            : IPUPAJHUAA
  Serial number        : NS1250G0115
  Manufacture date     : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state  : up
  Operational state    : unprovisioned
  Time of last boot    : 2017/10/23 12:27:29
  Current alarm state  : alarm active
  Firmware revision status : acceptable
-----
Chassis Control Module (CCM) Information

CCM Slot               : A
Equipped               : yes
Hardware Data
  Part number          : 82-0263-03
  CLEI code            :
  Serial number        : AQUA03-30
  Manufacture date     :
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number :
  Administrative state  : up
  Operational state    : up
  Temperature          : 45C
  Temperature threshold : 75C
  Time of last boot    : N/A
  Current alarm state  : alarm cleared
Hardware Resources (Power-Zone 1)
  Voltage
    Minimum            : 51.67 Volts (10/23/2017 12:31:59)
    Current             : 51.67 Volts
    Peak                : 52.68 Volts (10/23/2017 12:29:52)
  Wattage
    Minimum             : 21.29 Watts (10/23/2017 12:30:20)

```

```

Current                : 21.98 Watts
Peak                  : 22.66 Watts (10/23/2017 12:29:50)
Max Required          : 22.00 Watts
Amperage
Minimum              : 0.41 Amps (10/23/2017 12:30:20)
Current              : 0.43 Amps
Peak                 : 0.43 Amps (10/23/2017 12:29:50)

CCM Slot              : B
Equipped              : yes
Hardware Data
Part number           : 82-0263-03
CLEI code             :
Serial number         : AQUA03-10
Manufacture date      :
Manufacturing deviations : (Not Specified)
Manufacturing assembly number :
Administrative state   : up
Operational state     : up
Temperature           : 44C
Temperature threshold : 75C
Time of last boot     : N/A
Current alarm state    : alarm cleared
Hardware Resources (Power-Zone 1)
Voltage
Minimum              : 51.63 Volts (10/23/2017 12:32:00)
Current              : 51.63 Volts
Peak                 : 52.59 Volts (10/23/2017 12:30:14)
Wattage
Minimum              : 22.66 Watts (10/23/2017 12:30:40)
Current              : 23.35 Watts
Peak                 : 23.35 Watts (10/23/2017 12:30:14)
Max Required          : 22.00 Watts
Amperage
Minimum              : 0.45 Amps (10/23/2017 12:30:14)
Current              : 0.46 Amps
Peak                 : 0.46 Amps (10/23/2017 12:31:32)
=====
*A:PE-1#

```

Sample Output: show chassis (showing 7750 SR with Ethernet satellites configured)

```

*A:PE-1# show chassis
=====
System Information
=====
Name                : PE-1
Type                : 7750 SR-12
Chassis Topology    : Standalone
Location            : (Not Specified)
Coordinates         : (Not Specified)
CLLI code           :
Number of slots     : 12
Oper number of slots : 12
Num of faceplate ports/connectors : 165
Num of physical ports : 165
Critical LED state   : Off
Major LED state      : Off

```

```

Minor LED state           : Off
Over Temperature state    : OK
Base MAC address          : 10:e8:78:b8:21:47
Admin chassis mode        : d
Oper chassis mode         : d
Fabric Speed              : 10 Gig
FP Generations             : FP3
System Profile             : none

```

```

=====
Chassis Summary
=====

```

```

Chassis   Role           Status
-----

```

```

1          Standalone     up
esat-1     Ethernet Satellite up

```

```

=====
*A:PE-1#

```

Output Fields: show chassis

[Table 50](#) describes the output fields for the **show chassis** command.

Table 50 **Output Fields: show chassis**

Label	Description
Name	The system name for the router.
Type	Displays the router model number.
Chassis Topology	Indicates the inter-chassis topology mode in which the system is operating. Standalone indicates that the system is comprised of a single physical router chassis. Extended (XRS-40) on a 7950 XRS based system indicates that two router chassis are connected together in a “back-to-back” topology with no additional switch fabric chassis. An extended chassis topology is comprised of two XRS-20 chassis and is also known as an XRS-40 system.
Chassis role	Specifies the Chassis Roles. Standalone — Specifies the value for all non-7950 XRS SR OS systems and for 7950 XRS-20 standalone systems. 7950 XRS-40 Master 7950 XRS-40 Extension
Location	The system location for the device.
Coordinates	A user-configurable string that indicates the Global Positioning System (GPS) coordinates for the location of the chassis. For example: • N 45 58 23, W 34 56 12 • N37 37' 00 latitude, W122 22' 00 longitude • N36*39.246' W121*40.121'

Table 50 **Output Fields: show chassis (Continued)**

Label	Description (Continued)
CLLI Code	The Common Language Location Identifier (CLLI) that uniquely identifies the geographic location of places and certain functional categories of equipment unique to the telecommunications industry.
Number of slots	The number of slots in the 7450 ESS and 7750 SR chassis that are available for plug-in cards. The total number includes the IOM/CCM slot(s) and the CPM/CFM slots.
Number of ports	The total number of ports currently installed in this chassis. This count does not include the Ethernet ports on the CPMs/CFMs/CCMs that are used for management access.
Oper Number of Slots	The number of slots in the 7950 XRS chassis that are available for XCM cards and CPM cards operating as the active or standby CPM. Oper Number of Slots goes to 24 when Chassis Topology is Extended.
Critical LED state	The current state of the Critical LED in this chassis.
Major LED state	The current state of the Major LED in this chassis.
Minor LED state	The current state of the Minor LED in this chassis.
Base MAC address	The base chassis Ethernet MAC address.
FP Generations	Identifies the FP generations for cards in the system.
System Profile	The scaling profile in use by the system.
Over Temperature state	Indicates if there is currently an over temperature condition (OK = not currently over temp).
Admin chassis mode	The configured chassis mode (chassis mode D).
Oper chassis mode	The current chassis mode.
Part number	The part number of the particular hardware assembly. In the show chassis output, the first set of Hardware Data output is for the chassis mid-plane.
CLEI code	The Common Language Equipment Code of the particular hardware assembly.
Serial number	The serial number of the particular hardware assembly.
Manufacture date	The manufacture date of the particular hardware assembly.
Manufacturing string	The factory inputted manufacturing text string for the particular hardware assembly.
Manufacturing deviations	Additional manufacturing data.
Manufacturing assembly number	Additional manufacturing data.
Time of last boot	The date and time the most recent boot occurred.

Table 50 **Output Fields: show chassis (Continued)**

Label	Description (Continued)
Current alarm state	Displays the alarm conditions for the specific board.
Number of fan trays	The total number of fan trays installed in this chassis.
Number of fans	The total number of fans installed in this chassis.
Fan tray number	The ID for each fan tray installed in the chassis
Fan tray type	Indicates the type of fan tray: Chassis - standard fan tray PCM - dedicated to cooling PCMs in the 7950 XRS-20e Universal chassis
Operational status	Current status of the fan tray.
Speed	Indicates the speed of the fans as a percentage.
Status	Current status of the particular hardware assembly (up, down, or not equipped).
Number of power supplies	The number of power supplies installed in the chassis.
Power supply number	The ID for each power supply installed in the chassis.
Power supply type	The basic type of the power supply.
Power supply model	The model of the power supply.
CCM Slot	The identifier of the CCM (A or B).
Equipped	Indicates if the CCM is detected as physically present.
Temperature	The current temperature detected by the particular hardware assembly.
Temperature threshold	The temperature at which the particular hardware assembly considers an over temperature condition to exist.
CCM	Number of Chassis Control Modules on this unit (for the 7750 SR-c12 and 7750 SR-c4 only).
Part number	The CCM part number.
CLEI code	The code used to identify the router.
Serial number	The CCM serial number. Not user modifiable.
Manufacture date	The chassis manufacture date. Not user modifiable.
Manufacturing string	Factory-inputted manufacturing text string. Not user modifiable.
Administrative state	Up — The card is administratively up. Down — The card is administratively down.

Table 50 **Output Fields: show chassis (Continued)**

Label	Description (Continued)
Operational state	Up — The card is operationally up. Down — The card is operationally down.
Temperature	The internal chassis temperature.
Temperature threshold	The value above which the internal temperature must rise in order to indicate that the temperature is critical.
Time of last boot	The date and time the most recent boot occurred.
Current alarm state	Displays the alarm conditions for the CCM.

datapath

Syntax	datapath slot[/mda[/connector[/port]]] [mac mac-chip-number] [fp fp-number] [detail]
Context	show
Description	This command displays the data path mapping for the specified card, MDA, FP, MAC, connector, and port combinations. Fully qualified port identifiers are displayed for ports when the detail or port parameters are specified. This command is supported on FP4-based hardware only.
Parameters	<i>slot</i> — Specifies the slot number of the card in the chassis to display. The maximum slot number is platform dependent. See the <i>Installation Guides</i> for more information. Values 1 to 10 <i>mda</i> — Specifies the MDA number in the slot to display. This is an integer value relative to the card. Values 1 to 2 <i>connector</i> — Specifies the connector number to display. This is an integer value relative to the MDA. Values c1 to c36 <i>port</i> — Specifies the connector port number to display. The full port identifier is displayed. This is an integer value relative to the connector. Values 1 to 10 <i>fp-number</i> — Specifies the FP number to display. This is an integer value relative to the card. Values 1 to 8

mac-chip-number — Specifies the MAC number to display. This is an integer value relative to the MDA.

Values 1 to 12

detail — Displays the full port identifiers for all ports with configured connectors.

Output The following output is an example of the **show datapath** command.

Sample Output:

```
*A:PE# show datapath 1/2/c1/1

=====
Card      MDA      FP      MAC Chip Num  Connector  Ports
-----
1          2          1          1             c1         1/2/c1/1
=====

*A:PE# show datapath 1/2 fp 1

=====
Card      MDA      FP      MAC Chip Num  Connector  Ports
-----
1          2          1          1             c1
1          2          1          1             c2
1          2          1          1             c3
1          2          1          1             c4
1          2          1          1             c5
1          2          1          1             c6
=====

*A:PE# show datapath 1/2 mac 1 detail

=====
Card      MDA      FP      MAC Chip Num  Connector  Ports
-----
1          2          1          1             c1         1/2/c1/1
1          2          1          1             c2         1/2/c2/1
1          2          1          1             c3         1/2/c3/1
1          2          1          1             c4         1/2/c4/1
1          2          1          1             c5         1/2/c5/1
1          2          1          1             c6         1/2/c6/1
=====

*A:PE#
```

power-management

Syntax **power-management [zone]**

Context show>chassis

Description This command shows the power management requirement and utilization information for the 7950 XRS.

Parameters *zone* — The chassis power zone.

Values 1, 2

Output The following output is an example of power management information, and [Table 51](#) describes the output fields.

Sample Output: show chassis power-management

```
*A:Dut-A# show chassis power-management
```

```
=====
Chassis Information
```

```
=====
Power Management Information
```

```
Power Management Mode           : basic
Power Safety Level               : 100%
Power Safety Alert               : 0 watts
Power-Zone                      : 1
Number of PEQs                  : 12

PEQ number                      : 1
  PEQ Equipped Type             : apeq-dc-2000
  PEQ Provisioned Type          : apeq-dc-2000
  Status                        : shutdown
  Input Feed Status             : input B down
  Hardware Data
    Part number                 : 3HE07114AARA01
    CLEI code                   : IPUPAJHUAA
    Serial number               : NS1250G0116
    Manufacture date            : 12202012
    Manufacturing deviations     : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state          : down
  Operational state             : down
  Time of last boot             : 2014/01/07 11:01:44
  Current alarm state           : alarm active
```

```
PEQ number                      : 2
  PEQ Equipped Type             : apeq-dc-2000
  PEQ Provisioned Type          : (Not Specified)
  Status                        : up
  Input Feed Status             : input B down
  Hardware Data
    Part number                 : 3HE07114AARA01
    CLEI code                   : IPUPAJHUAA
    Serial number               : NS1249G0022
    Manufacture date            : 12202012
    Manufacturing deviations     : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state          : up
  Operational state             : unprovisioned
  Time of last boot             : 2014/01/07 11:01:44
  Current alarm state           : alarm active
```

```
PEQ number                      : 3
  PEQ Equipped Type             : apeq-dc-2000
```

```

PEQ Provisioned Type      : apeq-dc-2000
Status                    : up
Input Feed Status         : input B down
Hardware Data
  Part number             : 3HE07114AARA01
  CLEI code                : IPUPAJHUAA
  Serial number            : NS1250G0141
  Manufacture date         : 12202012
  Manufacturing deviations  : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state     : up
  Operational state        : up
  Time of last boot        : 2014/01/07 11:01:44
  Current alarm state      : alarm active

PEQ number                : 4
PEQ Equipped Type         : apeq-dc-2000
PEQ Provisioned Type      : apeq-dc-2000
Status                    : up
Input Feed Status         : input B down
Hardware Data
  Part number             : 3HE07114AARA01
  CLEI code                : IPUPAJHUAA
  Serial number            : NS1249G0201
  Manufacture date         : 12202012
  Manufacturing deviations  : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state     : up
  Operational state        : up
  Time of last boot        : 2014/01/07 11:01:44
  Current alarm state      : alarm active

PEQ number                : 5
PEQ Equipped Type         : apeq-dc-2000
PEQ Provisioned Type      : apeq-dc-2000
Status                    : up
Input Feed Status         : input B down
Hardware Data
  Part number             : 3HE07114AARA01
  CLEI code                : IPUPAJHUAA
  Serial number            : NS1250G0123
  Manufacture date         : 12202012
  Manufacturing deviations  : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state     : up
  Operational state        : up
  Time of last boot        : 2014/01/07 11:01:44
  Current alarm state      : alarm active

PEQ number                : 6
PEQ Equipped Type         : apeq-dc-2000
PEQ Provisioned Type      : apeq-dc-2000
Status                    : up
Input Feed Status         : input B down
Hardware Data
  Part number             : 3HE07114AARA01
  CLEI code                : IPUPAJHUAA
  Serial number            : NS1250G0061
  Manufacture date         : 12182012

```

```

Manufacturing deviations      : (Not Specified)
Manufacturing assembly number : 8205320107
Administrative state          : up
Operational state             : up
Time of last boot             : 2014/01/07 11:01:44
Current alarm state           : alarm active

PEQ number                    : 7
PEQ Equipped Type             : apeq-dc-2000
PEQ Provisioned Type          : apeq-dc-2000
Status                        : up
Input Feed Status             : input B down
Hardware Data
  Part number                 : 3HE07114AARB01
  CLEI code                   : IPUPAJHUAA
  Serial number               : NS13226A310
  Manufacture date            : 06042013
  Manufacturing deviations     : (Not Specified)
  Manufacturing assembly number : 82-0532-02
  Administrative state        : up
  Operational state           : up
  Time of last boot           : 2014/01/07 11:01:44
  Current alarm state         : alarm active

PEQ number                    : 8
PEQ Equipped Type             : apeq-dc-2000
PEQ Provisioned Type          : apeq-dc-2000
Status                        : up
Input Feed Status             : input B down
Hardware Data
  Part number                 : 3HE07114AARA01
  CLEI code                   : IPUPAJHUAA
  Serial number               : NS1250G0152
  Manufacture date            : 12202012
  Manufacturing deviations     : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state        : up
  Operational state           : up
  Time of last boot           : 2014/01/07 11:01:44
  Current alarm state         : alarm active

PEQ number                    : 9
PEQ Equipped Type             : apeq-dc-2000
PEQ Provisioned Type          : apeq-dc-2000
Status                        : up
Input Feed Status             : input B down
Hardware Data
  Part number                 : 3HE07114AARA01
  CLEI code                   : IPUPAJHUAA
  Serial number               : NS1250G0122
  Manufacture date            : 12202012
  Manufacturing deviations     : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state        : up
  Operational state           : up
  Time of last boot           : 2014/01/07 11:01:44
  Current alarm state         : alarm active

PEQ number                    : 10

```

```

PEQ Equipped Type           : apeq-dc-2000
PEQ Provisioned Type        : apeq-dc-2000
Status                       : up
Input Feed Status           : input B down
Hardware Data
  Part number                : 3HE07114AARA01
  CLEI code                  : IPUPAJHUAA
  Serial number              : NS1250G0146
  Manufacture date           : 12202012
  Manufacturing deviations    : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state        : up
  Operational state           : up
  Time of last boot           : 2014/01/07 11:01:44
  Current alarm state         : alarm active

PEQ number                   : 11
PEQ Equipped Type           : apeq-dc-2000
PEQ Provisioned Type        : apeq-dc-2000
Status                       : up
Input Feed Status           : input B down
Hardware Data
  Part number                : 3HE07114AARA01
  CLEI code                  : IPUPAJHUAA
  Serial number              : NS1249G0202
  Manufacture date           : 12202012
  Manufacturing deviations    : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state        : up
  Operational state           : up
  Time of last boot           : 2014/01/07 11:01:44
  Current alarm state         : alarm active

PEQ number                   : 12
PEQ Equipped Type           : apeq-dc-2000
PEQ Provisioned Type        : apeq-dc-2000
Status                       : up
Input Feed Status           : input B down
Hardware Data
  Part number                : 3HE07114AARA01
  CLEI code                  : IPUPAJHUAA
  Serial number              : NS1250G0115
  Manufacture date           : 12202012
  Manufacturing deviations    : (Not Specified)
  Manufacturing assembly number : 8205320107
  Administrative state        : up
  Operational state           : up
  Time of last boot           : 2014/01/07 11:01:44
  Current alarm state         : alarm active
=====

PCM number                   : 1
  PCM Equipped Type          : not detected
  PCM Provisioned Type        : (not provisioned)
PCM number                   : 2
  PCM Equipped Type          : not detected)
  PCM Provisioned Type        : (not provisioned)

```

```

PCM number                : 3
  PCM Equipped Type        : not detected
  PCM Provisioned Type      : (not provisioned)
PCM number                : 4
  PCM Equipped Type        : not detected
  PCM Provisioned Type      : (Not provisioned)
PCM number                : 5
  PCM Equipped Type        : not detected
  PCM Provisioned Type      : (Not provisioned)
PCM number                : 6
  PCM Equipped Type        : indeterminate
  PCM Provisioned Type      : (Not provisioned)
PCM number                : 7
  PCM Equipped Type        : indeterminate
  PCM Provisioned Type      : (Not provisioned)
PCM number                : 8
  PCM Equipped Type        : indeterminate
  PCM Provisioned Type      : (Not provisioned)
PCM number                : 9
  PCM Equipped Type        : indeterminate
  PCM Provisioned Type      : dual-pcm
PCM number                : 10
  PCM Equipped Type        : quad-pcm
  PCM Provisioned Type      : (Not provisioned)
Power-Zone                : 1
Status                    : up
Input Feed Status         : power to all inputs

Hardware Data
  Part number              : Sim Part#
  CLEI code                : Sim CLEI
  Serial number            : peq-12
  Manufacture date         : 01012003
  Manufacturing deviations  : Sim MfgDeviation peq-12
  Manufacturing assembly number : 01-2345-67

PCM number                : 11
  PCM Equipped Type        : quad-pcm
  PCM Provisioned Type      : quad-pcm
Power-Zone                : 1
Status                    : up
Input Feed Status         : power to all inputs
Hardware Data
  Part number              : Sim Part#
  CLEI code                : Sim CLEI
  Serial number            : peq-12
  Manufacture date         : 01012003
  Manufacturing deviations  : Sim MfgDeviation peq-12
  Manufacturing assembly number : 01-2345-67

PCM number                : 12
  PCM Equipped Type        : quad-pcm
  PCM Provisioned Type      : quad-pcm
Power-Zone                : 1
Status                    : up
Input Feed Status         : power to all inputs
Hardware Data
  Part number              : Sim Part#

```

```
CLEI code           : Sim CLEI
Serial number       : peq-12
Manufacture date    : 01012003
Manufacturing deviations : Sim MfgDeviation peq-12
Manufacturing assembly number : 01-2345-67
```

Chassis Control Module (CCM) Information

```
CCM Slot           : A
Equipped           : yes
Hardware Data
  Part number      : Sim Part#
  CLEI code        : Sim CLEI
  Serial number    : ccm-1
  Manufacture date : 01012003
  Manufacturing deviations : Sim MfgDeviation ccm-1
  Manufacturing assembly number : 01-2345-67
  Administrative state : up
  Operational state  : up
  Temperature       : -128C
  Temperature threshold : 75C
  Time of last boot : N/A
  Current alarm state : alarm cleared
Hardware Resources (Power-Zone 1)
  Voltage
    Minimum        : 52.80 Volts (09/10/2014 18:09:58)
    Current         : 52.80 Volts
    Peak           : 52.80 Volts (09/10/2014 18:09:58)
  Wattage
    Minimum        : 20.00 Watts (09/10/2014 18:09:58)
    Current         : 20.00 Watts
    Peak           : 20.00 Watts (09/10/2014 18:09:58)
    Max Required   : 22.00 Watts
  Amperage
    Minimum        : 0.38 Amps (09/10/2014 18:09:58)
    Current         : 0.38 Amps
    Peak           : 0.38 Amps (09/10/2014 18:09:58)
```

```
CCM Slot           : B
Equipped           : yes
Hardware Data
  Part number      :
  CLEI code        :
  Serial number    :
  Manufacture date :
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number :
  Administrative state : up
  Operational state  : up
  Temperature       : 0C
  Temperature threshold : 75C
  Time of last boot : N/A
  Current alarm state : alarm cleared
Hardware Resources (Power-Zone 1)
  Voltage
    Minimum        : 0.00 Volts (N/A)
    Current         : 0.00 Volts
```

```

Peak                               :      0.00 Volts  (N/A)
Wattage
Minimum                            :      0.00 Watts  (N/A)
Current                            :      0.00 Watts
Peak                               :      0.00 Watts  (N/A)
Max Required                        :      0.00 Watts
Amperage
Minimum                            :      0.00 Amps   (N/A)
Current                            :      0.00 Amps
Peak                               :      0.00 Amps   (N/A)
=====
Chassis 2 Detail
=====

```

Table 51 **Output Fields: show chassis power-management**

Label	Description
Power Management Mode	Specifies the configured power management mode: None, Basic, or Advanced.
Power Safety Level	Specifies the configured Power Safety Level, which is a percentage of the worst case power consumption level.
Power Safety Alert	Specifies the configured power level in watts, which causes the system to raise an alarm if the available power level drops below a set level.
Power-Zone	Specifies the chassis power zone.
Number of PEQs	Specifies the total number of APEQs installed.
PEQ number:	Specifies the APEQ to which the information is associated
PEQ Equipped Type	Specifies the APEQ type installed.
PEQ Provisioned Type	Specifies the APEQ type provisioned.
Status	Specifies the APEQ status.
Input Feed Status	Specifies the feed status. This field indicates that power is detected on all inputs or lists the inputs where no power is detected.
PCM Equipped Type	Indicates the type of PCM installed in the specified PIM or PCM slot: • (empty slot) • indeterminate • quad-pcm Some PCMs can only be verified for presence based on power available at the associated APEQ. These are indicated as indeterminate. When the equipped type is quad-pcm, the additional fields of Power-Zone, Status, and Input Feed Status and Hardware Data are displayed.

Table 51 **Output Fields: show chassis power-management (Continued)**

Label	Description
PCM Provisioned Type	One of the following types: • (not specified) • dual-pcm • quad-pcm For a given PCM number, if the equipped type is not detected and the provisioned type is not specified, the PCM is not displayed.
Hardware Data:	
Part number	The APEQ part number.
CLEI code	The APEQ CLEI code.
Serial number	The APEQ serial number.
Manufacture date	The date the APEQ was manufactured.
Manufacturing deviations	Specifies any manufacturing deviations.
Manufacturing assembly number	The APEQ assembly number.
Administrative state	Specifies the administrative state of the APEQ.
Operational state	Specifies the operational state of the APEQ.
Time of last boot	Indicates the time stamp of the last system restart.
Current alarm state	Indicates the current alarm state.

requirements

Syntax	requirements [detail]
Context	show>chassis>power-management
Description	This command displays maximum power requirements for the installed devices.
Parameters	detail — Displays detailed information.
Output	The following output is an example of power management requirements information, and Table 52 describes the output fields.

Sample Output: show chassis power-management requirements

```
*A:Dut-A#    show chassis power-management requirements
```

```
=====
```



```

Chassis Power Requirements
=====
                SUPPLY                                REQUIREMENTS
Power Capacity   : 22000.00 Watts                    Fan           : 1800.00 Watts ( 8%)
Safety Level     : 13203.00 Watts (100%)              IO Module      : 1395.00 Watts ( 6%)
Alert Level      :      0.00 Watts                    CPM Module     : 408.00 Watts ( 2%)
                                                Fabric Module  : 1600.00 Watts ( 7%)
                                                MDA Module    : 7956.00 Watts ( 36%)
                                                CCM Module    : 44.00 Watts ( 0%)
                                                Total Required: 13203.00 Watts ( 60%)
=====

```

Table 52 **Output Fields: show chassis power-management requirements**

Label	Description
SUPPLY	
Power Capacity	Indicates the total amount of power available to the chassis.
Safety Level	Specifies the configured Power Safety Level, which is a percentage of the worst case power consumption level.
Alert Level	Specifies the configured power level in watts, which causes the system to raise an alarm if the available power level drops below a set level.
REQUIREMENTS	
Fan	Specifies the amount of power required for each fan tray.
IO Module	Specifies the amount of power required for each IO Module.
CPM Module	Specifies the amount of power required for each CPM.
Fabric Module	Specifies the amount of power required for each SFM.
MDA Module	Specifies the amount of power required for each line card.
Total Required	Specifies the total amount of power required for all system elements.

utilization

- Syntax** **utilization [detail]**
- Context** show>chassis>power-management
- Description** This command displays power management utilization.
- Parameters** **detail** — Displays detailed breakdown of devices.

Output The following output is an example of power management utilization information, and [Table 53](#) describes the output fields.

Sample Output: show chassis power-management utilization

```
*A:Dut-A# show chassis power-management utilization
=====
Chassis Power Utilization
=====
                SUPPLY                PEAK DEMAND
Power Capacity  : 22000.00 Watts      Fan       : 695.00 Watts ( 3%)
Safety Level    : 13203.00 Watts (100%) IO Module  : 7163.09 Watts ( 33%)
Alert Level     : 0.00 Watts          CPM Module : 392.86 Watts ( 2%)
                                     Fabric Module: 1622.25 Watts ( 7%)
                                     MDA Module  : 6023.07 Watts ( 27%)+
                                     CCM Module  : 50.82 Watts ( 0%)
                                     Peak Util.   : 9924.02 Watts ( 45%)
                                     CURRENT DEMAND
                                     Current Util. : 9623.01 Watts ( 44%)
=====
+ Power utilization of device already included in IO Module value
=====
```

Sample Output: show chassis power-management utilization detail

```
*A:Dut-A# show chassis power-management utilization detail
=====
Chassis Power Utilization (detail)
=====
                SUPPLY                PEAK DEMAND
Power Capacity  : 22000.00 Watts      Fan
Power Supply 1  : 0.00 Watts          Fan 1      : 324.00 Watts ( 1%)
Power Supply 2  : 2000.00 Watts       Fan 2      : 371.00 Watts ( 2%)
Power Supply 3  : 2000.00 Watts       Fan 3      : 372.00 Watts ( 2%)
Power Supply 4  : 2000.00 Watts       PCM Fan 1  : 50.00 Watts ( 2%)
Power Supply 5  : 2000.00 Watts       PCM Fan 2  : 50.00 Watts ( 2%)
Power Supply 6  : 2000.00 Watts      IO Module
Power Supply 7  : 2000.00 Watts      Slot 1     : 812.19 Watts ( 4%)
Power Supply 8  : 2000.00 Watts      Slot 2     : 784.18 Watts ( 4%)
Power Supply 9  : 2000.00 Watts      Slot 3     : 799.01 Watts ( 4%)

Safety Level    : 13203.00 Watts (100%) CPM Module
Alert Level     : 0.00 Watts          Slot A     : 197.12 Watts ( 1%)
                                     Slot B     : 195.74 Watts ( 1%)
                                     Fabric Module
                                     Sfm 1     : 201.92 Watts ( 1%)
                                     Sfm 2     : 203.30 Watts ( 1%)
                                     Sfm 3     : 205.36 Watts ( 1%)
                                     Sfm 4     : 201.92 Watts ( 1%)
                                     Sfm 5     : 201.24 Watts ( 1%)
                                     Sfm 6     : 203.98 Watts ( 1%)
                                     Sfm 7     : 202.61 Watts ( 1%)
                                     Sfm 8     : 201.92 Watts ( 1%)
                                     MDA Module
                                     MDA 1/1    : 342.86 Watts ( 2%)+
                                     MDA 1/2    : 334.06 Watts ( 2%)+
                                     MDA 2/1    : 330.77 Watts ( 2%)+
```

```

MDA 2/2      :   331.87 Watts ( 2%)+
MDA 3/1      :   335.16 Watts ( 2%)+
MDA 3/2      :   343.96 Watts ( 2%)+
MDA 5/1      :   331.87 Watts ( 2%)+
MDA 5/2      :   326.37 Watts ( 1%)+
MDA 6/1      :   336.26 Watts ( 2%)+
MDA 6/2      :   332.97 Watts ( 2%)+
MDA 7/1      :   339.56 Watts ( 2%)+
MDA 7/2      :   332.97 Watts ( 2%)+
MDA 8/1      :   339.56 Watts ( 2%)+
MDA 8/2      :   328.57 Watts ( 1%)+
MDA 9/1      :   336.26 Watts ( 2%)+
MDA 9/2      :   331.87 Watts ( 2%)+
MDA 10/1     :   336.26 Watts ( 2%)+
MDA 10/2     :   332.97 Watts ( 2%)+
CCM Module
  CCM 1      :    24.73 Watts ( 0%)
  CCM 2      :    26.10 Watts ( 0%)
Peak Util.   :  9924.02 Watts ( 45%)
CURRENT DEMAND
Current Util. :  9613.68 Watts ( 44%)
=====
+ Power utilization of device already included in IO Module value
=====

```

Table 53 **Output Fields: show chassis power-management utilization**

Label	Description
SUPPLY	
Power Capacity	Indicates the total amount of power available to the chassis.
Safety Level	Specifies the configured Power Safety Level, which is a percentage of the worst case power consumption level.
Alert Level	Specifies the configured power level in watts, which causes the system to raise an alarm if the available power level drops below a set level.
DEMAND	
Fan	Specifies the amount of power utilized for the fan tray indicated.
IO Module	Specifies the amount of power utilized for the IO Module indicated.
CPM Module	Specifies the amount of power utilized for the CPM indicated.
Fabric Module	Specifies the amount of power utilized for the SFM indicated.
MDA Module	Specifies the amount of power utilized for the line card indicated.
Current Util.	Specifies the total amount of power utilized for all system elements.
Peak Util.	Specifies peak utilization starting from boot up.

card

Syntax

```

card slot-number cpu [sample-period seconds]
card slot-number fp fp-number dist-cpu-protection
card slot-number fp fp-number policy-accounting
card slot-number memory-pools
card state
card [slot-number]
card [slot-number] detail
card slot-number fp fp-number fwd-engine drop-reason statistics
card slot-number fp fp-number ingress queue-group mode {access | network}
card slot-number [detail] fp fp-number ingress queue-group queue-group-name instance
    [instance-number] mode {access | network} [statistics]

```

Context show

Description This command shows card information.

If no command line parameters are specified, a card summary for all cards is displayed.

Parameters *slot-number* — Displays information for the specified card slot.

Values Depending on the chassis model, IOM/XCM slots are numbered from 1 to 10.
SF/CPM slots are A, B. The 7950 XRS has additional slots C, D (upper or lowercase).
SFM slots are not addressed as cards. See the **show sfm** command.

cpu — Displays CPU utilization.

seconds — Displays the number of seconds over which to sample CPU task utilization.

Values 1 to 300

fp-number — Displays information for the specified forwarding plane.

Values 1 to 8

dist-cpu-protection — Displays the distributed CPU protection information.

policy-accounting — Displays the policy accounting information.

memory-pools — Displays the memory pools for the card.

state — Displays provisioned and equipped card and MDA information.

detail — Displays detailed card information.

fwd-engine — Displays forwarding engine information.

drop-reason — Displays drop reason information.

statistics — Displays statistics information.

ingress — Specifies statistics are for an ingress queue group.

- queue-group** — Displays queue group information.
- mode** — Specifies the mode in which the card will operate.
- access** — Displays information for the queue groups related to access mode.
- network** — Displays information for the queue groups related to network mode.

Output See the following sections for sample outputs:

- [Sample Output: show card \(showing the 7750 SR-12\)](#)
- [Sample Output: show card \(showing the 7950 XRS\)](#)
- [Output Fields: show card](#)
- [Sample Output: show card state \(showing 7750 SR-12 Chassis\)](#)
- [Sample Output: show card state \(showing 7750 SR-c12 Chassis\)](#)
- [Sample Output: show card state \(showing 7450 ESS-12 Chassis\)](#)
- [Sample Output: show card state \(showing a 7950 XRS Chassis\)](#)
- [Output Fields: show card state](#)
- [Sample Output: show card <slot-number> detail \(showing IOM3-XP Card\)](#)
- [Sample Output: show card <slot-number> detail \(showing XCM-20 Card\)](#)
- [Output Fields: show card <slot-number> detail \(for an IOM or XCM Card\)](#)
- [Sample Output: show card <slot-number> detail \(showing SFM and SFM3-12 Cards\)](#)
- [Sample Output: show card <slot-number> detail \(showing CPM-x20 Cards\)](#)
- [Output Fields: show card <slot-number> detail \(for a SF/CPM\)](#)
- [Sample Output:](#)
- [Output Fields:](#)
- [Sample Output](#)
- [Output Fields:](#)

Sample Output: show card (showing the 7750 SR-12)

```
*A:ALU-99# show card
=====
Card Summary
=====
Slot      Provisioned Type      Admin Operational  Comments
          Equipped Type (if different)  State State
-----
1          iom3-xp                up    up
4          imm12-10gb-sf+        up    up
5          imm12-10gb-sf+        up    up
A          sfm4-12                up    up/active
B          sfm4-12                up    up/standby
```

Sample Output: show card (showing the 7950 XRS)

```
A:Dut-A# show card
=====
Card Summary
```

```
=====
Slot   Provisioned Type           Admin Operational  Comments
       Equipped Type (if different) State State
-----
1      xcm-x20                    up      up
2      xcm-x20                    up      up
A      cpm-x20                    up      up/active
B      cpm-x20                    up      up/standby
C      cpm-x20                    up      down/ext-stby
D      cpm-x20                    up      up/ext-actv
=====
```

Output Fields: show card

Table 54 describes the output fields for the **show card** command.

Table 54 Output Field: show card

Label	Description
Slot	The slot number of the card in the chassis.
Provisioned Type	The card type that is configured for the slot. Note: CPMs C and D do not appear in the summary unless the Chassis Topology is Extended (for the 7950 XRS-40).
Equipped Type	The card type that is actually populated in the slot. Note: CPMs C and D do not appear in the summary unless the Chassis Topology is Extended (for the 7950 XRS-40).
Admin State	Up — The card is administratively up. Down — The card is administratively down.
Operational State	Up — The card is operationally up. Down — The card is operationally down. active — The CPM is the Active CPM for the system (actively managing the system components, processing various protocols, etc). standby — The CPM is the Standby CPM. The standby is hot synchronized with the Active CPM. ext-actv — The CPM is operating in an Extension role in a 7950 XRS-40 system and is the active extension CPM for the chassis in which it sits. ext-stby — The CPM is operating in an Extension role in a 7950 XRS-40 system and is the standby extension CPM for the chassis in which it sits.

Sample Output: show card state (showing 7750 SR-12 Chassis)

```
A:ALA-48# show card state
=====
Card State
=====
```

Slot/ MDA	Provisioned Type	Equipped Type	Admin State	Operational State	Num Ports	Num MDA	Comments
1	iom3-xp	iom3-xp	up	up		2	
1/1	m60-10/100eth-tx	m60-10/100eth-tx	up	up	60		
1/2	m60-10/100eth-tx	m60-10/100eth-tx	up	up	60		
2	iom3-xp		up	provisioned		2	
3	iom3-xp		up	provisioned		2	
3/1	m12-chds3		down	provisioned	12		
4	iom3-xp		up	provisioned		2	
4/1	m12-chds3		up	provisioned	12		
4/2	m1-choc12-sfp		up	provisioned	1		
5	iom3-xp		up	provisioned		2	
5/2	m12-chds3		down	provisioned	12		
6	iom3-xp		up	provisioned		2	
6/1	m12-chds3		up	provisioned	12		
6/2	m1-choc12-sfp		up	provisioned	1		
7	iom3-xp		up	provisioned		2	
7/1	m12-chds3		up	provisioned	12		
7/2	m1-choc12-sfp		up	provisioned	1		
8	iom3-xp		up	provisioned		2	
8/2	m1-choc12-sfp		up	provisioned	1		
9	iom3-xp		up	provisioned		2	
10	iom3-xp		up	provisioned		2	
A	sfm3-12	sfm-400g	up	up			Active
B	sfm3-12		up	provisioned			Standby

=====

A:ALA-48#

Sample Output: show card state (showing 7750 SR-c12 Chassis)

```
A:7750-3>config>card# show card state
```

=====

Card State

=====

Slot/ Id	Provisioned Type	Equipped Type	Admin State	Operational State	Num Ports	Num MDA	Comments
1	iom-xp	iom-xp	up	up		12	
1/1	mcm-xp	mcm-xp	up	up			
1/3		mcm-xp	up	unprovisioned			
1/1	m60-10/100eth-tx	m60-10/100eth-tx	up	up		60	
1/5	c8-10/100eth-tx	c8-10/100eth-tx	up	up		8	
1/6		c1-1gb-sfp	up	unprovisioned			
1/7		c8-chds1	up	unprovisioned			
1/8		c4-ds3	up	unprovisioned			
1/9		c8-10/100eth-tx	up	unprovisioned			
1/10		c1-1gb-sfp	up	unprovisioned			
1/11		c8-chds1	up	unprovisioned			
1/12		c4-ds3	up	unprovisioned			
A	cfm-xp	cfm-xp	up	up			Active
B	cfm-xp		up	down			Standby

=====

A:7750-3>

Sample Output: show card state (showing 7450 ESS-12 Chassis)

```
A:ALA-42# show card state
=====
Card State
=====
Slot/   Provisioned   Equipped   Admin Operational   Num   Num   Comments
MDA     Type           Type              State  State   Ports  MDA
-----
1       iom3-xp        iom3-xp        up     up       2
1/1     m60-10/100eth-tx m60-10/100eth-tx up     up       60
1/2     m60-10/100eth-tx m60-10/100eth-tx up     up       60
2       iom3-xp        up            provisioned          2
2/1     m60-10/100eth-tx up            provisioned          60
2/2     m60-10/100eth-tx up            provisioned          60
3       iom3-xp        up            provisioned          2
4       iom3-xp        up            provisioned          2
5       iom3-xp        up            provisioned          2
5/1     m20-100eth-sfp up            provisioned          20
6       iom3-xp        up            provisioned          2
7       iom3-xp        up            provisioned          2
10      iom3-xp        up            provisioned          2
A       sfm3-7         sfm3-7         up     up
B       sfm3-12        up            provisioned
=====
A:ALA-42#
```

Sample Output: show card state (showing a 7950 XRS Chassis)

```
A:Dut-A# show card state
=====
Card State
=====
Slot/   Provisioned Type   Admin Operational   Num   Num   Comments
Id      Equipped Type (if different) State  State   Ports  MDA
-----
1       xcm-x20          up     up       2
1/1     cx20-10g-sfp     up     up       20
1/2     cx20-10g-sfp     up     up       20
2       xcm-x20          up     up       2
2/1     cx20-10g-sfp     up     up       20
A       cpm-x20          up     up
B       cpm-x20          up     up
=====
```

Output Fields: show card state

Table 55 describes the output fields for the **show card state** command.

Table 55 Output Fields: show card state

Label	Description
Slot/MDA	The slot number of the card in the chassis.
Provisioned Type	The card type that is configured for the slot.

Table 55 Output Fields: show card state (Continued)

Label	Description (Continued)
Equipped Type	The card type that is actually populated in the slot.
Admin State	Up The card is administratively up. Down — The card is administratively down.
Operational State	Up — The card is operationally up. Provisioned — There is no card in the slot but it has been pre-configured.
Num Ports	The number of ports available on the MDA.
Num MDA	The number of MDAs installed.
Comments	Indicates whether the SF/CPM is the active or standby.

Sample Output: show card <slot-number> detail (showing IOM3-XP Card)

```
A:Dut-A# show card 10 detail
=====
Card 10
=====
Slot          Provisioned      Equipped        Admin   Operational      Comments
              Card-type       Card-type       State   State
-----
10            iom3-xp         iom3-xp         up      up

IOM Card Specific Data
  Clock source           : none
  Named Pool Mode        : Disabled
  Fail On Error          : Disabled
  Available MDA slots     : 2
  Installed MDAs         : 1

FP 1 Specific Data
  WRED Admin State       : Out Of Service
  WRED buffer-allocation max : 2500
  WRED buffer-allocation min : 2500
  WRED reserved-cbs max    : 2500
  WRED reserved-cbs min    : 2500
  WRED Slope Policy       : default
  hi-bw-mc-srcEgress Alarm : 2
  hi-bw-mc-srcEgress Group : 0
  mc-path-mgmt Admin State : Out Of Service
  Ingress Bandwidth Policy : default
  Stable Pool Sizing       : False
  Ingress Buffer Allocation : 50.00
  Initial Extract Priority Mode : uniform
  HS Pool Policy           : None
  HS Fixed High Threshold Delta : default
  Generation              : FP3
```

```

Hardware Data
Platform type           : 7750
Part number             : 3HE03619AAAK01
CLEI code               : IPU3AC9EAA
Serial number           : NS1112F0955
Manufacture date        : 03182011
Manufacturing string    :
Manufacturing deviations :
Manufacturing assembly number : 82-0107-09
Administrative state    : up
Operational state       : up
Temperature              : 50C
Temperature threshold   : 75C
Software boot (rom) version : X-0.0.I3122 on Mon Oct 17 18:16:02 PDT 2011*
Software version        : TiMOS-I-8.0.B1-250 iom/hops Nokia 7750 SR*
Time of last boot       : 2011/11/15 08:44:52
Current alarm state     : alarm cleared
Base MAC address        : 8c:90:d3:a4:fb:33
Last bootup reason      : hard boot
Memory capacity         : 2,048 MB

```

Sample Output: show card <slot-number> detail (showing XCM-20 Card)

```

A:Dut-A# show card 1 detail
=====
Card 1
=====
Slot   Provisioned Type           Admin Operational   Comments
      Equipped Type (if different) State State
-----
1      xcm-x20                   up    up

IOM Card Specific Data
Clock source           : none
Named Pool Mode        : Disabled
Fail On Error          : Disabled
Available MDA slots    : 2
Installed MDAs         : 2

FP 1 Specific Data
WRED Admin State       : Out Of Service
WRED buffer-allocation max : 2500
WRED buffer-allocation min : 2500
WRED reserved-cbs max   : 2500
WRED reserved-cbs min   : 2500
WRED Slope Policy       : default
hi-bw-mc-srcEgress Alarm : 2
hi-bw-mc-srcEgress Group : 0
mc-path-mgmt Admin State : Out Of Service
Ingress Bandwidth Policy : default
Stable Pool Sizing      : False
Ingress Buffer Allocation : 50.00
Initial Extract Priority Mode : uniform
Generation              : N/A

FP 2 Specific Data
WRED Admin State       : Out Of Service

```

```

WRED buffer-allocation max      : 2500
WRED buffer-allocation min      : 2500
WRED reserved-cbs max           : 2500
WRED reserved-cbs min           : 2500
WRED Slope Policy                : default
hi-bw-mc-srcEgress Alarm        : 2
hi-bw-mc-srcEgress Group        : 0
mc-path-mgmt Admin State        : Out Of Service
Ingress Bandwidth Policy        : default
Stable Pool Sizing              : False
Ingress Buffer Allocation        : 50.00
Initial Extract Priority Mode    : uniform
Generation                      : FP3

Hardware Data
Platform type                   : 7950
Part number                     :
CLEI code                      :
Serial number                   : xx
Manufacture date               :
Manufacturing string            : (Not Specified)
Manufacturing deviations        : (Not Specified)
Manufacturing assembly number   : 82-0334-02
Administrative state            : up
Operational state               : up
Temperature                    : 45C
Temperature threshold           : 75C
Software boot (rom) version     : X-0.0.I3326 on Thu May 10 18:22:55 PDT
                                2012 by builder
Software version                : TiMOS-I-10.0.S209 iom/hops Nokia 7950 XRS
Time of last boot               : 2012/05/23 20:27:09
Current alarm state             : alarm cleared
Base MAC address                : 00:21:05:8a:ca:0b
Last bootup reason              : hard boot
Memory capacity                 : 4,096 MB
* indicates that the corresponding row element may have been truncated.
=====
A:Dut-A#

```

Output Fields: show card <slot-number> detail (for an IOM or XCM Card)

[Table 56](#) describes the output fields for the **show card detail** command.

Table 56 Output Fields: show card <slot-number> detail (for an IOM or XCM Card)

Label	Description
Clock source	Source of clock for the IOM. Currently this parameter always displays "none".
Available MDA slots	The number of MDA slots available on the IOM.
Installed MDAs	The number of MDAs installed on the IOM.

Table 56 **Output Fields: show card <slot-number> detail (for an IOM or XCM Card)**

Label	Description (Continued)
Initial Extract Priority Mode	The scheme used to select the initial drop priority of extracted control plane traffic. uniform — Initializes the drop priority of all extracted control traffic as high priority. l3-classify — Initializes the drop priority of all Layer 3 extracted control traffic (for example, BGP or OSPF) based on the QoS classification of the packets. Refer to <i>Classification-Based Priority for Extracted Protocol Traffic</i> in the 7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide for more details.
Part number	The IOM part number.
CLEI code	The Common Language Location Identifier (CLLI) code string for the router.
Serial number	The serial number. Not user-modifiable.
Manufacture date	The chassis manufacture date. Not user-modifiable.
Manufacturing string	Factory-inputted manufacturing text string. Not user-modifiable.
Manufacturing deviations	Displays a record of changes by manufacturing to the hardware or software that is outside the normal revision control process.
Administrative state	Up — The card is administratively up. Down — The card is administratively down.
Operational state	Up — The card is operationally up. Down — The card is operationally down.
Temperature	Internal chassis temperature.
Temperature threshold	The value above which the internal temperature must rise in order to indicate that the temperature is critical.
Software boot version	The version of the boot image.
Software version	The software version number.
Time of last boot	The date and time the most recent boot occurred.
Current alarm state	Displays the alarm conditions for the specific board.
Base MAC address	Displays the base MAC address of the hardware component.
Memory capacity	Displays the memory capacity of the card.
Generation	Displays the FP generation of the assembly. N/A is displayed if there is no XMA installed in the card slot.

Sample Output: show card <slot-number> detail (showing SFM and SFM3-12 Cards)

```
B:Dut-D# show card A detail
```

```
=====
```

```

Card A
=====
Slot      Provisioned      Equipped      Admin      Operational
Card-type Card-type      Card-type      State      State
-----
A          sfm-400g          sfm-200g          up          up/standby

BOF last modified      : N/A
Config file version    :
Config file last modified : N/A
Config file last saved  : N/A
CPM redundancy status   : standby ready

Flash - cf1:
  Administrative State : up
  Operational state    : not equipped

Flash - cf2:
  Administrative State : up
  Operational state    : not equipped

Flash - cf3:
  Administrative State : up
  Operational state    : up
  Serial number        : 109917C1204W513
  Firmware revision    : HDX 2.1
  Model number         : SanDisk SDCFBI-128
  Size                 : 125,038 KB
  Free space           : 116,238 KB

Hardware Data
  Part number          : 3HE00316AAAA01
  CLEI code            : IPUCACMFAA
  Serial number        : NS043610538
  Manufacture date     : 02252005
  Manufacturing string  :
  Manufacturing deviations :
  Administrative state  : up
  Operational state     : up
  Temperature          : 43C
  Temperature threshold : 75C
  Software boot version : X-2.0.R12 on Mon Jan 24 12:03:16 PST 2007 *
  Software version      : TiMOS-C-4.0.private cpm/hops Nokia 7750 SR
  Time of last boot     : 2007/04/11 09:39:23
  Current alarm state   : alarm cleared
  Base MAC address     : 00:03:fa:30:7c:3c
  Memory capacity       : 2,016 MB
=====
B:Dut-D#
B:NS082761964# show card B detail
=====
Card B
=====
Slot      Provisioned      Equipped      Admin      Operational      Comments
Card-type Card-type      Card-type      State      State
-----
B          sfm3-12          sfm3-12          up          up/active
BOF last modified      : N/A

```

```

Config file version      : WED AUG 11 19:33:06 2010 UTC
Config file last modified : N/A
Config file last saved   : N/A
M/S clocking ref state   : primary

Flash - cf1:
  Administrative State    : up
  Operational state       : not equipped

Flash - cf2:
  Administrative State    : up
  Operational state       : not equipped

Flash - cf3:
  Administrative State    : up
  Operational state       : up
  Serial number           : 365ST295S3453SC01311
  Firmware revision       : V2.23
  Model number            : SILICONSYSTEMS INC 256MB
  Size                    : 253,932 KB
  Free space              : 121,368 KB

Hardware Data
  Platform type           : 7750
  Part number             : 3HE03617AAAA01
  CLEI code               : IPUCAN4FAA
  Serial number           : NS987456321
  Manufacture date        : 05072010
  Manufacturing string     :
  Manufacturing deviations :
  Manufacturing assembly number :
  Administrative state     : up
  Operational state       : up
  Temperature             : 34C
  Temperature threshold   : 75C
  Software boot (rom) version : X-0.0.I2627 on Thu Jun 10 18:03:16 PDT 2010*
  Software version         : TiMOS-C-0.0.private cpm/hops Nokia 7750 SR
  Time of last boot       : 2010/08/24 13:07:56
  Current alarm state     : alarm cleared
  Base MAC address        : 00:03:fa:1b:d7:16
  Memory capacity         : 4,096 MB
  System timing oscillator type : OCXO
=====

```

Sample Output: show card <slot-number> detail (showing CPM-x20 Cards)

```

*A:bksim3107# show card A detail
=====
Card A
=====
Slot Provisioned Type Admin Operational Comments
Equipped Type (if different) State State
-----
A cpm-x20 up up/active

BOF last modified : 2013/05/15 12:33:22
Config file version : FRI MAR 08 13:24:58 2013 UTC
Config file last modified : 2013/05/15 12:34:22

```

```
Config file last saved : 2013/05/15 12:36:22
M/S clocking ref state : primary

Flash - cf1:
Administrative State : up
Operational state : up
Serial number : serial-1
Firmware revision : v1.0
Model number : PC HD 1
Size : 1,950 MB
Free space : 1,950 MB

Flash - cf2:
Administrative State : up
Operational state : up
Serial number : serial-2
Firmware revision : v1.0
Model number : PC HD 2
Size : 0 Bytes
Free space : 0 Bytes

Flash - cf3:
Administrative State : up
Operational state : up
Serial number : serial-3
Firmware revision : v1.0
Model number : PC HD 3
Size : 18,432 Bytes
Free space : 6,144 Bytes
Hardware Data
Platform type : 7950
Part number : Sim Part#
CLEI code : Sim CLEI
Serial number : card-11
Manufacture date : 01012003
Manufacturing deviations : Sim MfgDeviation card-11
Manufacturing assembly number : 01-2345-67
Administrative state : up
Operational state : up
Temperature : -1C
Temperature threshold : 75C
Software boot (rom) version : simulated
Software version : TiMOS-C-11.0.R2 cpm/i386 Nokia 7950 XRS*
Time of last boot : 2013/05/13 08:10:33
Current alarm state : alarm cleared
Base MAC address : ac:9f:0b:00:00:01
Memory capacity : 3,072 MB

Inter Chassis CPM Interconnect
CPM Interconnect Port 1
Oper State : up
SFF Status : operational

CPM Interconnect Port 2
...
* indicates that the corresponding row element may have been truncated.
=====

*A:Dut-A# show card D detail
```

```

=====
Card D
=====
Slot   Provisioned Type           Admin Operational   Comments
       Equipped Type (if different) State State
-----
D      cpm-x20                   up    up/ext-stby

```

```

BOF last modified      : N/A
Config file version    :
Config file last modified : N/A
Config file last saved  : N/A
M/S clocking ref state : secondary

```

```

Flash - cf1:
  Administrative State : up
  Operational state    : up
  Serial number        : WE11K6300191
  Firmware revision    : 2.1ME
  Model number         : WDC SSD-D0128S-7117
  Size                 : 122,089 MB
  Free space           : 122,089 MB

```

```

Flash - cf2:
  Administrative State : up
  Operational state    : not equipped

```

```

Flash - cf3:
  Administrative State : up
  Operational state    : up
  Serial number        : SPG2012061404165
  Firmware revision    : 20101222
  Model number         : SMART CF
  Size                 : 3,907 MB
  Free space           : 3,802 MB

```

```

Hardware Data
  Platform type        : 7950
  Part number          : 3HE07116AARB01
  CLEI code            : IPUCA9T1AA
  Serial number        : NS13426D067
  Manufacture date     : 03162014
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number : 82-0488-05
  Administrative state  : up
  Operational state     : up
  Temperature          : 39C
  Temperature threshold : 75C
  Software boot (rom) version : X-12.0.B1-120 on Wed Jul 16 18:55:26 PDT
                          2014 by builder
  Software version      : TiMOS-C-12.0.B1-120 cpm/hops64 Nokia_7950 XRS
                          7950 Copyright (c) 2000-2016 Nokia
                          All rights reserved. All use subject to
                          applicable license agreements.
                          Built on Wed Jul 16 19:26:12 PDT 2014 by
                          builder in /rel12.0/b1/B1-120/panos/main
  Time of last boot     : 2014/07/17 13:41:28

```



```

Current alarm state      : alarm cleared
Base MAC address        : 00:d0:f6:f3:3c:9e
Memory capacity         : 8,192 MB

Hardware Resources (Power-Zone 2)
Voltage
  Minimum               : 53.10 Volts (07/17/2014 12:40:28)
  Current               : 53.16 Volts
  Peak                  : 54.15 Volts (07/17/2014 12:18:27)
Wattage
  Minimum               : 151.10 Watts (07/17/2014 13:31:23)
  Current               : 202.61 Watts
  Peak                  : 208.79 Watts (07/17/2014 13:00:07)
  Max Required          : 204.00 Watts
Amperage
  Minimum               : 3.69 Amps (07/17/2014 12:18:27)
  Current               : 3.82 Amps
  Peak                  : 3.94 Amps (07/17/2014 13:07:25)

Inter Chassis CPM Interconnect
CPM Interconnect Port 1
  Oper State            : up
  SFF Status            : operational
CPM Interconnect Port 2
  Oper State            : up
  SFF Status            : operational
=====

```

Output Fields: show card <slot-number> detail (for a SF/CPM)

[Table 57](#) describes the output fields for the **show card detail** command for a SF/CPM card.

Table 57 Output Fields: show card <slot-number> detail (for a SF/CPM)

Label	Description
Slot	The slot of the card in the chassis.
Card Provisioned	The SF/CPM type that is configured for the slot. Note: CPMs C and D will not show up in the summary unless the Chassis Topology is Extended (XRS-40).
Card Equipped	The SF/CPM type that is actually populated in the slot. Note: CPMs C and D will not show up in the summary unless the Chassis Topology is Extended (XRS-40).
Admin State	Up — The SF/CPM is administratively up. Down — The SF/CPM is administratively down.
Operational State	Up — The SF/CPM is operationally up. Down — The SF/CPM is operationally down.
Inter chassis cpm interconnect	Up — The CPM is operationally up. Down — The CPM is operationally down.

Table 57 **Output Fields: show card <slot-number> detail (for a SF/CPM) (Continued)**

Label	Description (Continued)
BOF last modified	The date and time of the most recent BOF modification.
Config file version	The configuration file version.
Config file last modified	The date and time of the most recent config file modification.
Config file last modified	The date and time of the most recent config file modification.
Config file last saved	The date and time of the most recent config file save.
CPM card status	active — The card is acting as the primary (active) CPM in a redundant system. standby — The card is acting as the standby (secondary) CPM in a redundant system.
Administrative state	Up — The CPM is administratively up. Down — The CPM is administratively down.
Operational state	Up — The CPM is operationally up. Down — The CPM is operationally down.
Serial number	The compact flash part number. Not user modifiable.
Firmware revision	The firmware version. Not user modifiable.
Model number	The compact flash model number. Not user modifiable.
Size	The amount of space available on the compact flash card.
Free space	The amount of space remaining on the compact flash card.
Part number	The SF/CPM part number.
CLEI code	The code used to identify the router.
Serial number	The SF/CPM part number. Not user modifiable.
Manufacture date	The chassis manufacture date. Not user modifiable.
Manufacturing string	Factory-set manufacturing text string. Not user modifiable.
Administrative state	Up — The card is administratively up. Down — The card is administratively down.
Operational state	Up — The card is operationally up. Down — The card is operationally down.
Time of last boot	The date and time the most recent boot occurred.
Current alarm state	Displays the alarm conditions for the specific board.

Table 57 **Output Fields: show card <slot-number> detail (for a SF/CPM) (Continued)**

Label	Description (Continued)
Status	Displays the current status.
Temperature	Internal chassis temperature.
Temperature threshold	The value above which the internal temperature must rise in order to indicate that the temperature is critical.
Software boot version	The version of the boot image.
Memory capacity	The total amount of memory.

Sample Output:

show card slot-number detail (showing IOM with license level)

```
*A:bkvm18# show card 1 detail
=====
Card 1
=====
Slot          Provisioned Type          Admin Operational   Comments
              Equipped Type (if different)  State State
-----
1             iom5-e:hel200g+                up    up
IOM Card Licensing Data
  Licensed Level           : hel200g+
  Description               : 1.2T /w Agg to 2.4T, High Edge Routing
  Licensed Bandwidth Used   : 900 Gbps
IOM Card Specific Data
  Clock source              : none
  Named Pool Mode           : Disabled
  Fail On Error             : Disabled
  Reset On Recoverable Error : Disabled
  Available MDA slots       : 2
  Installed MDAs            : 0
FP 1 Specific Data
  WRED Admin State          : Out Of Service
  WRED buffer-allocation max : 2500
  WRED buffer-allocation min : 2500
  WRED reserved-cbs max     : 2500
  WRED reserved-cbs min     : 2500
  WRED Slope Policy         : default
  hi-bw-mc-srcEgress Alarm  : disabled
  hi-bw-mc-srcEgress Group  : 0
  mc-path-mgmt Admin State  : Out Of Service
  Ingress Bandwidth Policy  : default
  Stable Pool Sizing        : False
  Ingress Buffer Allocation  : 50.00
  Initial Extract Priority Mode : uniform
  HS Pool Policy            : None
  HS Fixed High Threshold Delta : default
  Generation                : FP4
Hardware Data
  Platform type             : N/A
  Part number               :
```

```

CLEI code           :
Serial number       :
Manufacture date    :
Manufacturing deviations : (Not Specified)
Manufacturing assembly number :
Administrative state : up
Operational state   : provisioned
Software boot (rom) version : (Not Specified)
Software version    : (Not Specified)
Time of last boot   : N/A
Current alarm state : alarm cleared
Base MAC address    :
Last bootup reason  :
Memory capacity     : 0 MB

```

Output Fields:

show card slot-number detail (showing IOM with license level)

[Table 58](#) describes the output fields for the **show card detail** command with IOM licensing information.

Table 58 **Output Fields: show card slot-number detail (showing IOM with license level)**

Label	Description
Slot	The slot of the card in the chassis.
Provisioned Type	The IOM type and level that is configured for the slot.
Equipped Type (if different)	The IOM type and level that is actually populated in the slot if it is different than the Provisioned Type.
Admin State	Up — The IOM is administratively up. Down — The IOM is administratively down.
Operational State	Up — The IOM is operationally up. Down — The IOM is operationally down.
Licensed Level	License level of the card.
Description	Description of the license.
Licensed Bandwidth Used	The total connector bandwidth configured on the assembly.

Sample Output

show card slot-number fp 1 fwd-engine drop-reason statistics

```
*A:cses-V93# show card 1 fp 1 fwd-engine drop-reason statistics
```

```
=====
```

```
Card 1 FP 1 Egress Forwarding Engine Drop Reason Statistics
=====
Egress Encapsulation Exceeded          0
MTU Exceeded                          0
=====
```

Output Fields:

show card slot-number fp 1 fwd-engine drop-reason statistics

[Table 59](#) describes the output fields for the **show card fp 1 fwd-engine drop-reason** command.

Table 59 Output Fields: show card slot-number fp 1 fwd-engine drop-reason

Label	Description
Card	The card in the chassis.
FP	The egress forwarding plane hardware component.
Egress Encapsulation Exceeded	Packet size exceeded the egress encapsulation size.
MTU Exceeded	Packet size exceeded the egress MTU.

cflowd

Syntax	cflowd
Context	show
Description	This command displays cflowd information.

elmi

Syntax	elmi
Context	show
Description	This command displays Ethernet Link Management Interface (eLMI) information.

[Table 60](#) describes the output fields for the **show elmi** command.

Table 60 **Output Fields: show elmi**

Field	Description
Link Status	Status of the E-LMI protocol when the elmi mode is set to uni-n. Link Status will indicate up if eLMI mode is set to "none".
T391	Pooling timer used by UNI-C. UNI-N will send the consecutive single EVC asynchronous status messages every (T391/10) rounded to the second interval.
T392	Pooling verification timer for UNI-N
N393	Status counter for UNI-N.
Rx Enq. Time	Last time a status inquiry message was received from UNI-C.
Rx Enq Msg	Number of status inquiry messages received.
Rx Check Time	Last time a status inquiry E-LMI check message was received.
Rx Inv. SeqNum	Counts the number of E-LMI messages received with invalid sequence number.
Enq Timeouts	Counts the number of T392 timer expired.
Tx Status Time	Last time a status message was sent by UNI-N.
Tx Status Msg	Number of status messages sent by UNI-N.
Tx Check Time	Last time a status eLMI check message was sent by UNI-N.
Tx Async Status Msg	Counter for single EVC asynchronous status messages sent by UNI-N.
Discard Msg	Counter for the status inquiry messages discarded due to errors.

evc

Syntax	evc [<i>port-id</i> [vlan <i>vlan-id</i>]]
Context	show>elmi
Description	This command displays Ethernet Virtual Connections (EVC). No argument displays all the EVC on the service router. The port and VLAN arguments display information related to EVC associated with the port and VLAN.
Parameters	<i>port-id</i> — Displays information related to the EVCs configured on the port. Values <i>slot/mda/port</i> <i>vlan-id</i> — Specifies the VLAN Identifier of the EVC. Values 0 to 4094, *
Output	The following output is an example of ELMi EVC information, and Table 61 describes the output fields.

Sample Output

```
*A:Dut-C# show elmi evc
=====
ELMI EVC Table
=====
Port      Vlan  Status  Type  Evc Id
-----
1/1/1     10    New-Act  P2p   EVC11110
1/1/3     30    New-Act  P2p   EVC11220
1/1/5     100   Act      P2p   EVC115100
1/1/5     200   Act      P2p   EVC115200
-----
Number of Evcs : 4
=====
*A:Dut-C#

A:Dut-C# show elmi evc 1/1/5
=====
ELMI EVC Table
=====
Port      Vlan  Status  Type  Evc Id
-----
1/1/5     100   Act      P2p   EVC115100
1/1/5     200   Act      P2p   EVC115200
-----
Number of Evcs : 2
=====
A:Dut-C#

*A:Dut-C# show elmi evc 1/1/5 vlan 100
=====
Evc Detailed Information
=====
Port      : 1/1/5          vlanId      : 100
Evc Status : Act          Evc Type    : P2p
Evc Identifier: EVC115100
=====
*A:Dut-C#
```

Table 61 **Output Fields: show elmi evc**

Field	Description
Port	Port identifier.
Vlan	Vlan identifier.
Status	New-Act — New and active. Act — Active.
Type	Connection type.
EVC Id	EVC identifier.

uni

Syntax	uni [<i>port-id</i>]
Context	show>elmi
Description	This command displays information about ELMI mode, status, number of EVCs (SAPs) configured on the port for all of the ports on the service router.
Parameters	<i>port-id</i> — Displays UNI information for the specified port. Values <i>slot/mda/port</i>
Output	The following output is an example of elmi uni information.

Sample Output: show elmi uni

```
*A:Dut-C# show elmi uni
=====
ELMI UNI-N Table
=====
Port      Mode   Status  #Evcs  Uni Identifier
-----
1/1/1     None   Up       0      10/100 Ethernet TX
1/1/2     None   Up       0      port-21
1/1/3     None   Up       0      10/100 Ethernet TX
1/1/4     None   Up       0      10/100 Ethernet TX
1/1/5     Uni-N  Up       2      UNI115
1/1/6     None   Up       0      10/100 Ethernet TX
1/1/7     None   Up       0      10/100 Ethernet TX
1/1/8     None   Up       0      10/100 Ethernet TX
1/1/9     None   Up       0      10/100 Ethernet TX
1/1/10    None   Up       0      10/100 Ethernet TX
1/1/11    None   Up       0      10/100 Ethernet TX
1/1/12    None   Up       0      10/100 Ethernet TX
1/1/13    None   Up       0      10/100 Ethernet TX
1/1/14    None   Up       0      10/100 Ethernet TX
1/1/15    None   Up       0      10/100 Ethernet TX
1/1/16    None   Up       0      10/100 Ethernet TX
1/1/17    None   Up       0      10/100 Ethernet TX
...
=====
*A:Dut-C#

*A:Dut-C# show elmi uni 1/1/5
=====
Uni-N Detailed Information
=====
Uni Mode      : Uni-N                      Link Status      : Up
Uni Identifier: UNI115
T391          : 10 seconds                T392             : 15 seconds
N393          : 4                        UniType          : Bundling
Rx Enq. Time  : 02/18/2010 17:11:44      Tx Status Time   : 02/18/2010 17:11:44
Rx Enq Msg    : 24                       Tx Status Msg    : 24
Rx Check Time : 02/18/2010 17:12:34      Tx Check Time    : 02/18/2010 17:12:34
Rx Inv. SeqNum: 0                        Tx Async Status Msg : 0
```



```

Enq Timeouts : 0                               Discard Msg : 0
=====
*A:Dut-C#

```

eth-tunnel

Syntax	eth-tunnel {aps status} eth-tunnel tunnel-index [path path-index] [detail]
Context	show
Description	This command displays Ethernet tunnel information.
Parameters	aps — Shows APS Ethernet tunnel information. status — Shows Ethernet tunnel status information. tunnel-index — Specifies the tunnel index. Values 1 to 1024 path-index — Specifies the path index. Values 1 to 16 detail — Displays detailed information.
Output	The following output is an example of Ethernet tunnel information, and Table 62 describes the output fields.

Sample Output: show eth-tunnel

```

*A:PE-E# show eth-tunnel
=====
Ethernet Tunnel Groups
=====
Tunnel Admin Oper Protection Active Paths
ID State State Type 1 2 3 4 5 6 7 8
-----
1 Up Up g.8031-1to1 x 2
2 Up Up g.8031-1to1 1 x
=====
*A:PE-E#
*A:PE-E# show eth-tunnel aps
=====
Ethernet Tunnel APS Groups
=====
Tunnel Admin Oper Working Path Path Active Rx PDU
ID State State Protecting Path State Path Tx PDU
-----
1 Up Up 1 - 1/1/2 1 Down No BF010100 ( SF)
2 - 2/1/2 1 Up Yes BF010100 ( SF)
2 Up Up 1 - 2/1/2 2 Up Yes 0F000000 ( NR)
2 - 1/1/2 2 Down No EF000000 (SF-P)
=====
*A:PE-E#

```

```

*A:PE-E# show eth-tunnel 1
=====
Ethernet Tunnel Group 1 Information
=====
Description          : Eth Tunnel
IfIndex              : 1476395009
Admin State          : Up
Protection Type       : G.8031-1to1
MAC Address           : 00:1a:f0:44:d2:03
Oper State            : Up
Max Revert Time      : 1 seconds
Time to Revert       : N/A
Hold Down Time       : 0 centiseconds
-----

Ethernet Tunnel Group APS Information
-----
APS PDU Rx           : BF010100 ( SF)
APS PDU Tx           : BF010100 ( SF)
Switchover Time      : 05/28/2009 10:10:17
Defect Status        :
-----

Ethernet Tunnel Group Path Summary
-----
Path ID  Member  Control-Tag  Precedence  Admin/Oper  Active  Mgmt
-----
1        1/1/2    1            primary     Up/Down     No      Yes
2        2/1/2    1            secondary   Up/Up       Yes     No
=====

*A:PE-E#

*A:PE-E# show eth-tunnel 1 path 1
=====
Ethernet Tunnel Group 1 Path Information
=====
Description          : (Not Specified)
Member               : 1/1/2
Admin State          : Up
Control-Tag          : 1
Oper State           : Down
-----

Ethernet Tunnel Group Path APS Information
-----
Active Count         : 2
Active Time          : 0d 00:12:09
-----

Eth-Cfm Configuration Information
-----
Md-index             : 1
Ma-index             : 1
MepId                : 1
LowestDefectPri      : macRemErrXcon
Defect Flags         : bDefRemoteCCM
HighestDefect        : defRemoteCCM
Mac Address          : 00:16:4d:c0:c1:ca
ControlMep           : True
=====

*A:PE-E#

*A:PE-E# show eth-tunnel 1 path 1 detail
=====
Ethernet Tunnel Group 1 Detailed Path Information
=====
Description          : (Not Specified)
Member               : 1/1/2
Admin State          : Up
Control-Tag          : 1
Oper State           : Down
-----

```

```

Ethernet Tunnel Group Path APS Information
-----
Active Count      : 2                Active Time      : 0d 00:12:09
-----
Eth-Cfm Configuration Information
-----
Md-index          : 1                Direction        : Down
Ma-index          : 1                Admin            : Enabled
MepId             : 1                CCM-Enable       : Enabled
LowestDefectPri   : macRemErrXcon    HighestDefect     : defRemoteCCM
Defect Flags      : bDefRemoteCCM
Mac Address       : 00:16:4d:c0:c1:ca ControlMep       : True
CcmLtmPriority    : 7
CcmTx             : 0                CcmSequenceErr   : 0
Eth-Ais:          : Disabled
Eth-Tst:          : Disabled
LbRxReply         : 0                LbRxBadOrder     : 0
LbRxBadMsdu       : 0                LbTxReply        : 0
LbNextSequence    : 1                LtNextSequence    : 1
LtRxUnexplained   : 0
=====
*A:PE-E#

```

Table 62 **Output Fields: show eth-tunnel**

Field	Description
Tunnel Id	Numeric value from 1 to 64.
Admin State	Up — The eth tunnel is administratively up. Down — The eth tunnel is administratively down.
Oper State	Up — The eth tunnel is operationally up. Down — The eth tunnel is operationally down.
Protection Type	Two options: g8031-1to1 — Two members are allowed, but only one of them is active at one point in time. loadsharing — Multiple members can be active at one point in time.
Active Paths	Only two paths are supported.

fwd-path-ext

Syntax **fwd-path-ext [fpe fpe-id]**
fwd-path-ext fpe fpe-id associations

Context show

Description This command displays FPE information.

Parameters *fpe-id* — Specifies the FPE ID.

Values 1 to 64

associations — Displays a list of current fwd path extensions to which the FPE is associated.

Output The following output is an example of show fwd-path-ext command information, and [Table 64](#) describes the output fields.

Sample Output

```
A:CPM148>config>fwd-path-ext>fpe$ show fwd-path-ext fpe
=====
FPE Info
=====
FPE Id      Path Application  Description
          pxc/xc-a, xc-b
-----
1          pxc 1          pw-port      xyz
2          lag 1, lag 2    vxlan-term   abc
-----
Number of entries : 2
-----
SDP-Id Range: 1-100
=====
```

```
*A:CPM148>config>fwd-path-ext>fpe$ show fwd-path-ext fpe 1
=====
FPE Id: 2
=====
Description      : xyz
Path             : pxc 1
Pw Port         : Enabled      Oper: up
Vxlan Termination : Disabled
=====
```

```
*A:CPM148>config>fwd-path-ext>fpe$ show fwd-path-ext fpe 2
=====
FPE Id: 2
=====
Description      : abc
Path             : lag 1, lag 2
Pw Port         : Disabled
Vxlan Termination : Disabled      Oper: down & Can be down due to lag 1/
2 being down, PXC's within the lags being down etc
=====
```

```
*A:CPM148>config>fwd-path-ext>fpe$ show fwd-path-ext fpe 1 associations
=====
pw-port associations
=====
```

```

Epipe Svc Id      PW-Port-Id
=====
110
220
330
=====
Vxlan-termination associations
=====
None
=====
*A:CPM148>config>fwd-path-ext>fpe$ show fwd-path-ext fpe 2 associations
=====
pw-port associations
=====
None
=====
Vxlan-termination associations
=====
tunnel-termination
=====
10.1.1.1
10.2.1.1
2001:db8::10
2001:db8::20
=====

```

Table 63 **Output Fields: show fwd-path-ext**

Field	Description
FPE Id	Displays the configured ID of the FPE.
Path pxc/ xc-a, xc-b	Displays the path associated with this FPE. This can be a single PXC (pair of PXC sub-ports .a and .b), or it can be a LAG with PXC sub-ports as member ports. PXC based LAG is used for redundancy and increased throughput of the FPE. In case of a PXC based LAG, the path is referred to as xc-a and xc-b.
Application	Displays the application associated with this FPE. Based on the application type, the system will internally configure logic on underlying PXC (or PXC based LAG) necessary to accommodate the application.
Description	Displays the user-configurable string used to describe the use of the FPE.
SDP-id Range	Displays the configurable SDP-id range used for internal SDPs that are utilized by FPEs.
Path	Path associated with this FPE. This can be a single PXC (pair of PXC sub-ports .a and .b), or it can be a LAG with PXC sub-ports as member ports. PXC based LAG is used for redundancy and increased throughput of the FPE.
PW-Port	Displays the FPE that is used by PW-port application.

Table 63 Output Fields: show fwd-path-ext (Continued)

Field	Description
VXLAN-Termination	Displays the FPE that is used by Vxlan-Termination application.
Oper	Displays the status of the application.
EPIPE svc Id	Displays the Epipe service ID with which the PW-port is associated.
PW-Port Id	Displays the PW-port ID that is associated with the EPIPE.
Tunnel termination	Displays the IP addresses on which the VXLAN tunnel is terminated.

interface-group-handler

Syntax	interface-group-handler [<i>index</i>]
Context	show
Description	This command displays Interface Group Handler (IGH) information. If no command line options are specified, a summary listing of all IGHs is displayed.
Parameters	<i>index</i> — Specifies the index. Values 1 to 100
Output	The following output is an example of Interface Group Handler information, and Table 64 describes the output fields.

Sample Output: show interface-group-handler

```

A:ALU-27# show interface-group-handler
=====
Interface Group Handler Summary Information
=====
IGH Index Admin      Number of Threshold
          State      Members
-----
1         Up         4         4
2         Up         2         2
=====
A:ALU-27#

A:ALU-27#show interface-group-handler 2
=====
Interface Group Handler 2 Information
=====
Admin Status      : Up
Threshold         : 2
Last Change       : 02/02/2010 18:10:04

```

```

-----
Interface Group Handler Protocol Information
-----
Protocol Oper Status  Active Links                                Up Time
-----
ipcp      up          2                                0d 00:15:04
ipv6cp    none         0                                N/A
mplscp    waiting      0                                N/A
osicp     none         0                                N/A
-----
Port 1/5/2.2 Information
-----
Protocol Oper Status                                Up Time
-----
ipcp      up          0d 00:15:05
ipv6cp    none         N/A
mplscp    running       N/A
osicp     none         N/A
-----
Port 1/5/2.3 Information
-----
Protocol Oper Status                                Up Time
-----
ipcp      up          0d 00:15:05
ipv6cp    none         N/A
mplscp    running       N/A
osicp     none         N/A
=====
A:ALU-27#

```

Table 64 Output Fields: show interface-group-handler

Field	Description
IGH Index	A value between 1 and 100 that identifies the specific interface group handler.
Admin State	Up — The interface group handler is administratively up. Down — The interface group handler administratively down.
Number of Members	Identifies the number of ports or channels in the group, up to a maximum of 8.
Threshold	Indicates the minimum number of active links that must be present for the interface group handler to be active.

mcm

Syntax **mcm** [*slot* [/*mcm*]] [**detail**]
Context show
Description This command displays MCM information.

If no command line options are specified, a summary output of all MCMs is displayed in table format.

Parameters *slot* — The slot number for which to display MCM information.

Values 1

mcm — The MCM number in the slot for which to display MCM information.

Values 7750 SR-c4 to 1, 3
 7750 SR-c12 to 1, 3, 5, 7, 9, 11

detail — Displays detailed MDA information.

Output The following output is an example of MCM information, and [Table 65](#) describes the output fields.

Sample Output: show mcm

```
A:7750-3# show mcm
=====
MCM Summary
=====
Slot Mcm   Provisioned      Equipped          Admin    Operational
      Mcm-type      Mcm-type          State      State
-----
1     1     mcm-xp           mcm-xp            up        up
      3     mcm-xp           mcm-xp            up        unprovisioned
=====
A:7750-3#

A:7750-3# show mcm 1
=====
MCM 1/1
=====
Slot Mcm   Provisioned      Equipped          Admin    Operational
      Mcm-type      Mcm-type          State      State
-----
1     1     mcm-xp           mcm-xp            up        up
=====
MCM 1/3
=====
Slot Mcm   Provisioned      Equipped          Admin    Operational
      Mcm-type      Mcm-type          State      State
-----
      3     mcm-xp           mcm-xp            up        unprovisioned
=====
A:7750-3#
```

Table 65 Output Fields: show mcm

Label	Description
Slot	The chassis slot number.
MCM	The MCM slot number.

Table 65 Output Fields: show mcm (Continued)

Label	Description
Provisioned MCM-type	The MCM type provisioned.
Equipped MCM-type	The MCM type installed.
Admin State	Up — Administratively up. Down — Administratively down.
Ops State	Up — Operationally up. Down — Operationally down.

mda

Syntax `mda slot [/mda] [detail]`

Context show

Description This command displays MDA\XMA information.

If no command line options are specified, a summary output of all MDAs is displayed in table format.

Parameters `slot` — Specifies the slot number for which to display MDA information.

Values 1 to 10

`mda` — Specifies the MDA number in the slot for which to display MDA information.

Values `slot [/mda]`

detail — Displays detailed MDA information.

Output See the following sections for output samples:

- [Sample Output: show mda \(showing 7450 ESS\)](#)
- [Sample Output: show mda \(showing 7950 XRS\)](#)
- [Output Fields: show mda](#)
- [Sample Output: show mda <slot/mda> detail](#)
- [Sample Output: show mda <slot/mda> detail \(showing CMA in Slot 1 on a 7750 SR-c12 or 7750 SR-c4 system\)](#)
- [Sample Output: show mda <slot/mda> detail \(showing Channelized MDA on a 7750 SR-7\)](#)
- [Sample Output: show mda <slot/mda> detail \(showing 7950 XRS\)](#)
- [Output Fields: show mda <slot/mda> detail](#)

Sample Output: show mda (showing 7450 ESS)

```

A:ALA-42# show mda
=====
MDA Summary
=====
Slot Mda   Provisioned      Equipped          Admin   Operational
      Mda-type      Mda-type          State    State
-----
1     1     m60-10/100eth-tx m60-10/100eth-tx up       up
      2     m60-10/100eth-tx m60-10/100eth-tx up       up
=====
A:ALA-42#

```

Sample Output: show mda (showing 7950 XRS)

```

A:Dut-A# show mda
=====
MDA Summary
=====
Slot  Mda   Provisioned Type   Equipped Type          Admin   Operational
              (if different)   State    State
-----
1     1     cx20-10g-sfp                up       up
      2     cx20-10g-sfp                up       up
2     1     cx20-10g-sfp                up       up
=====
A:Dut-A#

```

Output Fields: show mda

[Table 66](#) describes the output fields for the **show mda** command.

Table 66 Output Fields: show mda

Label	Description
Slot	The chassis slot number.
MDA	The MDA slot number.
Provisioned MDA-type	The MDA type provisioned.
Equipped MDA-type	The MDA type actually installed.
Admin State	Up — Administratively up. Down — Administratively down (e.g., shutdown).
Operational State	Up — Operationally up. Down — Operationally down.

Sample Output: show mda <slot/mda> detail

```

*A:Dut-A# show mda 5/1 detail

```

```

=====
MDA 5/1 detail
=====
Slot  Mda  Provisioned      Equipped      Admin  Operational
      Mda-type      Mda-type      State      State
-----
5      1      m20-1gb-xp-sfp      m20-1gb-xp-sfp      up      up

MDA Specific Data
Maximum port count      : 20
Number of ports equipped : 20
Network ingress queue policy : default
Capabilities      : Ethernet
Fail On Error      : disabled
Egress XPL error threshold : 1000
Egress XPL error window  : 60

Hardware Data
Platform type      : 7750
Part number      : 3HE03612AAAB01
CLEI code      : IPPAABFBAA
Serial number      : NS093464752
Manufacture date      : 08232009
Manufacturing string      :
Manufacturing deviations :
Manufacturing assembly number :
Administrative state      : up
Operational state      : up
Temperature      : 37C
Temperature threshold      : 75C
Software version      : N/A
Time of last boot      : 2011/11/15 11:32:49
Current alarm state      : alarm cleared
Base MAC address      : 00:23:3e:ea:38:4b

-----
QOS Settings
-----
Ing. Named Pool Policy      : None
Egr. Named Pool Policy      : None
=====

```

Sample Output: show mda <slot/mda> detail (showing CMA in Slot 1 on a 7750 SR-c12 or 7750 SR-c4 system)

```

A:7750-3# show mda 1/5 detail
=====
MDA 1/5 detail
=====
Slot  Mda  Provisioned      Equipped      Admin  Operational
      Mda-type      Mda-type      State      State
-----
1      5      c8-10/100eth-tx      c8-10/100eth-tx      up      up

MDA Specific Data
Maximum port count      : 8
Number of ports equipped : 8
Network ingress queue policy : default
Capabilities      : Ethernet
Fail On Error      : disabled

```

```
Egress XPL error threshold : 1000
Egress XPL error window   : 60
```

Hardware Data

```
Part number      : Sim Part#
CLEI code        : Sim CLEI
Serial number    : mda-5
Manufacture date : 01012003
Manufacturing string : Sim MfgString mda-5
Manufacturing deviations : Sim MfgDeviation mda-5
Administrative state : up
Operational state  : up
Temperature       : 33C
Temperature threshold : 75C
Time of last boot  : 2007/04/11 15:13:48
Current alarm state : alarm cleared
Base MAC address   : 04:7b:01:05:00:01
```

Sample Output: show mda <slot/mda> detail (showing Channelized MDA on a 7750 SR-7)

```
A:SR-7/Dut-C# show mda 5/1 detail
```

MDA 5/1 detail

```
=====
Slot Mda   Provisioned      Equipped      Admin   Operational
      Mda-type      Mda-type      State    State
-----
5    1    m1-choc12-sfp    m1-choc12-sfp    up      up
```

MDA Specific Data

```
Maximum port count      : 1
Number of ports equipped : 1
Transmit timing selected : CPM Card A
Sync interface timing status : Qualified
Network ingress queue policy : default
Capabilities             : Sonet, TDM, PPP, FR
Fail On Error           : disabled
Egress XPL error threshold : 1000
Egress XPL error window   : 60
```

```
Min channel size        : PDH DS0 Group
Max channel size         : PDH DS3
Max number of channels   : 512
Channels in use          : 0
```

Hardware Data

```
Part number      : 3HE00193AAAA01
CLEI code        :
Serial number    : NS042510655
Manufacture date : 07072004
Manufacturing string :
Manufacturing deviations :
Administrative state : up
Operational state  : up
```

```

Time of last boot          : 2007/04/11 12:51:48
Current alarm state       : alarm cleared
Base MAC address         : 00:03:fa:1a:7c:6f
=====
A:SR-7/Dut-C#

```

Sample Output: show mda <slot/mda> detail (showing 7950 XRS)

```

A:Dut-A# show mda 1/1 detail
=====
MDA 1/1 detail
=====
Slot  Mda   Provisioned Type  Equipped Type(if different) Admin   Operational
-----
1      1      cx20-10g-sfp                up      up

MDA Specific Data
Maximum port count       : 20
Number of ports equipped : 20
Network ingress queue policy : default
Capabilities              : Ethernet
Min channel size         : Sonet STS-192
Max channel size         : Sonet STS-192
Max number of channels    : 20
Channels in use          : 0

Hardware Data
Platform type            : 7950
Part number               :
CLEI code                :
Serial number            : GRA03-126
Manufacture date         :
Manufacturing string      : (Not Specified)
Manufacturing deviations  : (Not Specified)
Manufacturing assembly number : 82-0299-03
Administrative state      : up
Operational state        : up
Temperature               : 45C
Temperature threshold     : 75C
Software version          : N/A
Time of last boot        : 2012/05/23 20:30:55
Current alarm state      : alarm cleared
Base MAC address         : 8c:90:d3:be:69:8a
Firmware version         : I-10.0.S209

-----
QOS Settings
-----
=====
A:Dut-A#

```

Output Fields: show mda <slot/mda> detail

[Table 67](#) describes the output fields for the **show mda <slot/mda> detail** command.

Table 67 **Output Fields: show mda <slot/mda> detail**

Label	Description
Slot	The chassis slot number.
Mda	The MDA slot number.
Provisioned Mda-type	The provisioned MDA type.
Equipped Mda-type	The MDA type that is physically inserted into this slot in this chassis.
Admin State	Up — The MDA is administratively up. Down — The MDA is administratively down.
Operational State	Up — The MDA is operationally up. Down — The MDA is operationally down.
Failure Reason	This hardware component has failed.
Maximum port count	The maximum number of ports that can be equipped on the MDA card.
Number of ports equipped	The number of ports that are actually equipped on the MDA.
Transmit timing selected	Indicates the source for the timing used by the MDA.
Sync interface timing status	Indicates whether the MDA has qualified one of the timing signals from the CPMs.
Network Ingress Queue Policy	Specifies the network queue policy applied to the MDA to define the queueing structure for this object.
Capabilities	Specifies the minimum size of the port that can exist on the MDA.
Egress XPL error threshold	The Egress XPL Error Threshold value used by the fail-on-error feature.
Egress XPL error window	The Egress XPL Error Window value used by the fail-on-error feature.
Max channel size	Specifies the maximum size of the channel that can exist on the channelized MDA.
Channels in use	Applicable for SONET and TDM MDAs only. Indicates the total number of leaf SONET paths, TDM channels and bundles on the MDA which are presently provisioned for passing traffic.
Part number	The hardware part number.
CLEI code	The code used to identify the MDA.
Serial number	The MDA part number. Not user modifiable.
Manufacture date	The MDA manufacture date. Not user modifiable.
Manufacturing string	Factory-inputted manufacturing text string. Not user modifiable.

Table 67 **Output Fields: show mda <slot/mda> detail (Continued)**

Label	Description (Continued)
Administrative state	Up — The MDA is administratively up. Down — The MDA is administratively down.
Operational state	Up — The MDA is operationally up. Down — The MDA is operationally down.
Time of last boot	The date and time the most recent boot occurred.
Current alarm state	Displays the alarm conditions for the specific MDA.
Base MAC address	The base chassis Ethernet MAC address. Special purpose MAC addresses used by the system software are constructed as offsets from this base address.

qos

Syntax **qos {ingress | egress} buffer-allocation [detail]**
qos {ingress | egress} orphaned-queues

Context show>mda

Description This command shows MDA QoS information.

Parameters **ingress** — Displays ingress information.
egress — Displays egress information.
buffer-allocation — Shows named pool buffer allocation information.
detail — Shows detailed information.
orphaned-queues — Shows configured but not defined queues.

pools

Syntax **pools mda-id [/port] [access-app [pool-name | service service-id]] | queue-group queue-group-name [instance id]]]**
pools mda-id [/port] [network-app [pool-name | queue-group queue-group-name [instance id]]]
pools mda-id [/port] [direction [pool-name | service service-id] queue-group queue-group-name [instance id]]]

Context show

Description This command displays pool information.

Parameters **mda-id[/port]** — Displays the pool information of the specified MDA and port.

access-app — Specifies the pool application as either ingress or egress.

Values access-ingress, access-egress

instance-id — Specifies the identification of a specific instance of the queue-group.

Values 1 to 65535

pool-name — Displays the pool information of the specified QoS policy up to 32 characters.

Values access-ingress, access-egress

service-id — Displays pool information for the specified service.

Values 1 to 2148278381, *svc-name*: 64 char max

queue-group-name — Displays information for the specified queue group up to 32 characters.

network-app — Specifies the pool application as either ingress or egress.

Values network-ingress, network-egress

direction — Specifies the traffic direction.

Values ingress, egress

Output The following outputs are example of pool information, and [Table 68](#) describes the output fields.



Note: The pool shared in use stat only increases when a queue is asking for a buffer outside it's reserved size. If all the buffers in a pool are assigned to queues within their reserved size, then only the reserved in use size will increase. In case of resv CBS oversubscription (CBS sum for all queues is bigger than pool resvCbs), it is possible that pool resv in use stat can increase above the actual pool reserved size.

Output Sample: show pools

A:PE# show pools 5/1

=====					
Type	Id-FP	App.	Pool Name	Actual ResvCBS Admin ResvCBS	PoolSize

MDA	5/1-1	Acc-Ing	default	0	0
				30%	
MDA	5/1-1	Acc-Ing	MC Path Mgmt	18816	37632
				50%	
MDA	5/1-1	Acc-Egr	default	0	0
				30%	
MDA	5/1-1	Net-Ing	default	0	0
				Sum	
MDA	5/1-1	Net-Egr	default	0	0
				Sum	
Port	5/1/1	Acc-Ing	default	0	0
				30%	


```

Port      5/1/1      Acc-Egr default      0      0
                                     30%
Port      5/1/1      Net-Egr default     32256   75264
                                     Sum
Port      5/1/2      Acc-Ing default      0      0
                                     30%
Port      5/1/2      Acc-Egr default      0      0
                                     30%
Port      5/1/2      Net-Egr default     32256   75264
                                     Sum
Port      5/1/3      Acc-Ing default      0      0
                                     30%
Port      5/1/3      Acc-Egr default      0      0
                                     30%
Port      5/1/3      Net-Egr default     32256   75264
                                     Sum
Port      5/1/4      Acc-Ing default      0      0
                                     30%
Port      5/1/4      Acc-Egr default      0      0
                                     30%
Port      5/1/4      Net-Egr default     32256   75264
                                     Sum
Port      5/1/5      Acc-Ing default      0      0
                                     30%
Port      5/1/5      Acc-Egr default      0      0
                                     30%
Port      5/1/5      Net-Egr default     32256   75264
                                     Sum
=====
A:PE#

```

Output Sample: show pools network-egress

```

*A:PE# show pools 5/1/5 network-egress
=====
Pool Information
=====
Port           : 5/1/5
Application    : Net-Egr      Pool Name       : default
CLI Config. Resv CBS : Sum
Resv CBS Step  : 0%           Resv CBS Max    : 0%
Amber Alarm Threshold: 0%     Red Alarm Threshold : 0%
-----
Utilization      State      Start-Avg    Max-Avg      Max-Prob
-----
HiPlus-Slope     Down        85%          100%         80%
High-Slope       Down        70%          90%          80%
Low-Slope        Down        50%          75%          80%
Exceed-Slope     Down        30%          55%          80%
Time Avg Factor  : 7
Pool Total       : 75264 KB
Pool Shared      : 43008 KB      Pool Resv      : 32256 KB
-----
Current Resv CBS  Provisioned  Rising      Falling      Alarm
%age             all Queues  Alarm Thd   Alarm Thd    Color
-----
Sum              32232 KB   NA          NA           Green

```

```
Pool Total In Use      : 0 KB
Pool Shared In Use     : 0 KB
WA Shared In Use       : 0 KB
HiPlus-Slope Drop Pr* : 0
Lo-Slope Drop Prob     : 0
Pool Resv In Use       : 0 KB
Hi-Slope Drop Prob     : 0
Excd-Slope Drop Prob   : 0
=====
Queue : 1 Net=be Port=5/1/5
=====
FC Map      : be
Tap         : not-applicable
Admin PIR   : 10000000
Admin CIR   : 0
Admin FIR   : 0
Admin MBS   : 37632 KB
High-Plus Drop T* : 37632 KB
Low Drop Tail : 33792 KB
CBS         : 744 KB
Slope       : not-applicable
FP          : not-applicable
Oper PIR    : Max
Oper CIR    : 0
Oper FIR    : 0
Oper MBS    : 37632 KB
High Drop Tail : 37632 KB
Exceed Drop Tail : 33792 KB
Depth      : 0
=====
Queue : 2 Net=l2 Port=5/1/5
=====
FC Map      : l2
Tap         : not-applicable
Admin PIR   : 10000000
Admin CIR   : 2500000
Admin FIR   : 0
Admin MBS   : 37632 KB
High-Plus Drop T* : 37632 KB
Low Drop Tail : 33792 KB
CBS         : 2256 KB
Slope       : not-applicable
FP          : not-applicable
Oper PIR    : Max
Oper CIR    : 2500000
Oper FIR    : 0
Oper MBS    : 37632 KB
High Drop Tail : 37632 KB
Exceed Drop Tail : 33792 KB
Depth      : 0
=====
Queue : 3 Net=af Port=5/1/5
=====
FC Map      : af
Tap         : not-applicable
Admin PIR   : 10000000
Admin CIR   : 2500000
Admin FIR   : 0
Admin MBS   : 37632 KB
High-Plus Drop T* : 37632 KB
Low Drop Tail : 33792 KB
CBS         : 7488 KB
Slope       : not-applicable
FP          : not-applicable
Oper PIR    : Max
Oper CIR    : 2500000
Oper FIR    : 0
Oper MBS    : 37632 KB
High Drop Tail : 37632 KB
Exceed Drop Tail : 33792 KB
Depth      : 0
=====
Queue : 4 Net=l1 Port=5/1/5
=====
FC Map      : l1
Tap         : not-applicable
Admin PIR   : 10000000
Admin CIR   : 2500000
Admin FIR   : 0
Admin MBS   : 18816 KB
High-Plus Drop T* : 18816 KB
Low Drop Tail : 16896 KB
CBS         : 2256 KB
Slope       : not-applicable
FP          : not-applicable
Oper PIR    : Max
Oper CIR    : 2500000
Oper FIR    : 0
Oper MBS    : 18816 KB
High Drop Tail : 18816 KB
Exceed Drop Tail : 16896 KB
Depth      : 0
```

Slope : not-applicable

=====

Queue : 5 Net=h2 Port=5/1/5

=====

FC Map	: h2		
Tap	: not-applicable	FP	: not-applicable
Admin PIR	: 10000000	Oper PIR	: Max
Admin CIR	: 10000000	Oper CIR	: Max
Admin FIR	: 0	Oper FIR	: 0
Admin MBS	: 37632 KB	Oper MBS	: 37632 KB
High-Plus Drop T*	: 37632 KB	High Drop Tail	: 37632 KB
Low Drop Tail	: 33792 KB	Exceed Drop Tail	: 33792 KB
CBS	: 7488 KB	Depth	: 0
Slope	: not-applicable		

=====

Queue : 6 Net=ef Port=5/1/5

=====

FC Map	: ef		
Tap	: not-applicable	FP	: not-applicable
Admin PIR	: 10000000	Oper PIR	: Max
Admin CIR	: 10000000	Oper CIR	: Max
Admin FIR	: 0	Oper FIR	: 0
Admin MBS	: 37632 KB	Oper MBS	: 37632 KB
High-Plus Drop T*	: 37632 KB	High Drop Tail	: 37632 KB
Low Drop Tail	: 33792 KB	Exceed Drop Tail	: 33792 KB
CBS	: 7488 KB	Depth	: 0
Slope	: not-applicable		

=====

Queue : 7 Net=h1 Port=5/1/5

=====

FC Map	: h1		
Tap	: not-applicable	FP	: not-applicable
Admin PIR	: 10000000	Oper PIR	: Max
Admin CIR	: 10000000	Oper CIR	: 10000000
Admin FIR	: 0	Oper FIR	: 0
Admin MBS	: 18816 KB	Oper MBS	: 18816 KB
High-Plus Drop T*	: 18816 KB	High Drop Tail	: 18816 KB
Low Drop Tail	: 16896 KB	Exceed Drop Tail	: 16896 KB
CBS	: 2256 KB	Depth	: 0
Slope	: not-applicable		

=====

Queue : 8 Net=nc Port=5/1/5

=====

FC Map	: nc		
Tap	: not-applicable	FP	: not-applicable
Admin PIR	: 10000000	Oper PIR	: Max
Admin CIR	: 10000000	Oper CIR	: 10000000
Admin FIR	: 0	Oper FIR	: 0
Admin MBS	: 18816 KB	Oper MBS	: 18816 KB
High-Plus Drop T*	: 18816 KB	High Drop Tail	: 18816 KB
Low Drop Tail	: 16896 KB	Exceed Drop Tail	: 16896 KB
CBS	: 2256 KB	Depth	: 0
Slope	: not-applicable		

=====

* indicates that the corresponding row element may have been truncated.

*A:PE#

Output Sample: show pool access-ingress

*A:PE# show pools 5/1/1 access-ingress

```

=====
Pool Information
=====
Port                : 5/1/1
Application         : Acc-Ing          Pool Name           : default
CLI Config. Resv CBS : 30%(default)
Resv CBS Step       : 0%               Resv CBS Max         : 0%
Amber Alarm Threshold: 0%               Red Alarm Threshold  : 0%
-----
Utilization          State      Start-Avg    Max-Avg    Max-Prob
-----
HiPlus-Slope         Down        85%         100%       80%
High-Slope           Down        70%         90%        80%
Low-Slope            Down        50%         75%        80%
Exceed-Slope         Down        30%         55%        80%

Time Avg Factor      : 7
Pool Total           : 264192 KB
Pool Shared          : 184320 KB      Pool Resv            : 79872 KB
-----

-----
Current Resv CBS     Provisioned   Rising        Falling        Alarm
%age                 all Queues   Alarm Thd     Alarm Thd      Color
-----
30%                  0 KB        NA            NA             Green
Pool Total In Use    : 0 KB
Pool Shared In Use   : 0 KB      Pool Resv In Use   : 0 KB
WA Shared In Use     : 0 KB

HiPlus-Slope Drop Pr*: 0      Hi-Slope Drop Prob  : 0
Lo-Slope Drop Prob   : 0      Excd-Slope Drop Prob : 0

=====
Queue : 100->5/1/1->1
=====
FC Map              : be 12 af 11 h2 ef h1 nc
Tap                 : 5/1                FP                : 1
Admin PIR           : 10000000          Oper PIR          : Max
Admin CIR           : 0                 Oper CIR          : 0
Admin FIR           : 0                 Oper FIR          : 0
Admin MBS           : 12288 KB          Oper MBS          : 12288 KB
High-Plus Drop Tail : 12288 KB          High Drop Tail    : 12288 KB
Low Drop Tail       : 10944 KB          Exceed Drop Tail  : 9792 KB
CBS                 : 0 KB              Depth             : 0
Slope               : not-applicable

=====

=====
Queue : 100->5/1/1->1
=====

```

```

FC Map          : be l2 af l1 h2 ef h1 nc
Tap             : 1/*
Admin PIR       : 10000000
Admin CIR       : 0
Admin FIR       : 0
Admin MBS       : 12288 KB
High-Plus Drop Tail : 12288 KB
Low Drop Tail   : 10944 KB
CBS             : 0 KB
Slope          : not-applicable

=====

Queue : 100->5/1/1->11
=====
FC Map          : be l2 af l1 h2 ef h1 nc
Tap             : MCast
Admin PIR       : 10000000
Admin CIR       : 0
Admin FIR       : 0
Admin MBS       : 12288 KB
High-Plus Drop Tail : 12288 KB
Low Drop Tail   : 10944 KB
CBS             : 0 KB
Slope          : not-applicable

=====
* indicates that the corresponding row element may have been truncated.
*A:PE#

```

Sample Output: show pools access-egress

```

*A:PE# show pools 5/1/1 access-egress
=====
Pool Information
=====
Port          : 5/1/1
Application    : Acc-Egr
CLI Config. Resv CBS : 30%(default)
Resv CBS Step  : 0%
Amber Alarm Threshold: 0%
Pool Name      : default
Resv CBS Max   : 0%
Red Alarm Threshold : 0%
-----
Queue-Groups
-----
Queue-Group:Instance : policer-output-queues:1
-----
Utilization      State      Start-Avg   Max-Avg     Max-Prob
-----
HiPlus-Slope     Down      85%         100%        80%
High-Slope       Down      70%         90%         80%
Low-Slope        Down      50%         75%         80%
Exceed-Slope     Down      30%         55%         80%
Time Avg Factor  : 7
Pool Total       : 75264 KB
Pool Shared      : 52224 KB
Pool Resv        : 23040 KB
-----
Current Resv CBS  Provisioned  Rising      Falling     Alarm

```

```

%age          all Queues      Alarm Thd      Alarm Thd      Color
-----
30%           0 KB            NA             NA             Green
Pool Total In Use : 0 KB
Pool Shared In Use : 0 KB
WA Shared In Use   : 0 KB
HiPlus-Slope Drop Pr*: 0
Lo-Slope Drop Prob : 0
Hi-Slope Drop Prob : 0
Excd-Slope Drop Prob : 0
=====
Queue : accQGrp->policer-output-queues:1 (5/1/1) ->1
=====
FC Map          : not-applicable
Tap             : not-applicable
Admin PIR       : 10000000      Oper PIR        : Max
Admin CIR       : 0             Oper CIR        : 0
Admin MBS       : 12288 KB      Oper MBS        : 12288 KB
High-Plus Drop Tail : 12288 KB    High Drop Tail  : 12288 KB
Low Drop Tail   : 10944 KB      Exceed Drop Tail : 9792 KB
CBS             : 0 KB          Depth           : 0
Slope           : not-applicable
=====
Queue : accQGrp->policer-output-queues:1 (5/1/1) ->2
=====
FC Map          : not-applicable
Tap             : not-applicable
Admin PIR       : 10000000      Oper PIR        : Max
Admin CIR       : 0             Oper CIR        : 0
Admin MBS       : 12288 KB      Oper MBS        : 12288 KB
High-Plus Drop Tail : 12288 KB    High Drop Tail  : 12288 KB
Low Drop Tail   : 10944 KB      Exceed Drop Tail : 9792 KB
CBS             : 0 KB          Depth           : 0
Slope           : not-applicable
=====
Queue : 1->5/1/1->1
=====
FC Map          : be l2 af l1 h2 ef h1 nc
Tap             : not-applicable
Admin PIR       : 10000000      Oper PIR        : Max
Admin CIR       : 0             Oper CIR        : 0
Admin MBS       : 12288 KB      Oper MBS        : 12288 KB
High-Plus Drop Tail : 12288 KB    High Drop Tail  : 12288 KB
Low Drop Tail   : 10944 KB      Exceed Drop Tail : 9792 KB
CBS             : 0 KB          Depth           : 0
Slope           : not-applicable
=====
* indicates that the corresponding row element may have been truncated.
*A:PE#

```

Output Fields: show pools

[Table 68](#) describes the output fields for the **show pools** command.

Table 68 **Output Fields: show pools**

Label	Description
Type	Specifies the pool type.
ID-FP	Specifies the card-FP or MDA-FP or card, MDA, or port designation.
Application/Type	Specifies what the pool would be used for. The pools could be used for access or network traffic at either ingress or egress.
Pool Name	Specifies the name of the pool being used.
Resv CBS	Specifies the percentage of pool size reserved for CBS.
Utilization	Specifies the type of the slope policy.
State	The administrative status of the port.
Start-Avg	Specifies the percentage of the buffer utilized after which the drop probability starts to rise above 0.
Max-Avg	Specifies the percentage of the buffer utilized after which the drop probability is 100 percent. This implies that all packets beyond this point will be dropped.
Time Avg Factor	Specifies the time average factor the weighting between the previous shared buffer average utilization result and the new shared buffer utilization in determining the new shared buffer average utilization.
Actual ResvCBS	Specifies the actual percentage of pool size reserved for CBS.
Admin ResvCBS	Specifies the percentage of pool size reserved for CBS.
PoolSize	Specifies the size in percentage of buffer space. The value '-1' implies that the pool size should be computed as per fair weighting between all other pools.
Pool Total	Displays the total pool size.
Pool Shared	Displays the amount of the pool which is shared.
Pool Resv	Specifies the percentage of reserved pool size.
Pool Total In Use	Displays the total amount of the pool which is in use.
Pool Shared In Use	Displays the amount of the pool which is shared that is in use.

2.21.2.2 PEQ Show Commands

The PEQ commands apply to the 7950 XRS only.

pcm

- Syntax** `pcm [pcm-slot] [chassis chassis-id] [detail]`
- Context** show
- Description** This command displays PCM information.
- Parameters** *pcm-slot* — Specifies the PCM slot identifier.
Values 1 to 12
chassis-id — Specifies the chassis ID for the router chassis.
Values 1 to 2
detail — Displays detailed information.
- Output** The following output shows PCM information.

Sample Output: show pcm

```
*A:Dut-A# show pcm
=====
PCM Summary
=====
```

Chassis/ Slot	Provisioned Type Equipped Type (if diff)	Operational State	Input A B2 A2 B	Zone
1/9	(not provisioned) (present)	unprovisioned		1
1/10	(not provisioned) quad-pcm	unprovisioned	Y Y Y Y	1
1/11	quad-pcm	unprovisioned	Y Y Y Y	1
1/12	quad-pcm	unprovisioned	Y Y Y Y	1

peq

- Syntax** `peq [peq-slot] [chassis chassis-id] [detail]`
- Context** show
- Description** This command displays APEQ information.
- Parameters** *peq-slot* — Specifies the APEQ slot identifier.
Values 1 to 12
chassis-id — Specifies the chassis ID for the router chassis.
Values 1 to 2
detail — Displays detailed information.

Output The following outputs are examples of PEQ information, and [Table 69](#) describes the output fields.

Sample Output: show peq

```
*A:Dut-A# show peq
=====
PEQ Summary
=====
```

Chassis/ Slot	Provisioned Type Equipped Type (if diff)	Admin State	Operational State	Input A B	Zone	Input Mode
1/1	apeq-dc-2000	up	down	Y Y	1	N/A
1/2	(not provisioned) apeq-dc-2000	up	unprovisioned	Y Y	1	N/A
1/3	apeq-dc-2000	up	up	Y Y	1	N/A
1/4	apeq-dc-2000	up	up	Y Y	1	N/A
1/5	apeq-dc-2000	up	up	Y Y	1	N/A
1/6	apeq-dc-2000	up	down	Y Y	1	N/A
1/7	apeq-dc-2000	up	down	Y Y	1	N/A
1/8	apeq-dc-2000	up	down	Y Y	1	N/A
1/9	apeq-dc-2000	up	down	Y Y	1	N/A
1/10	apeq-dc-2000	up	down	Y Y	1	N/A
1/11	apeq-dc-2000	up	down	Y Y	1	N/A
1/12	apeq-dc-2000	up	down	Y Y	1	N/A
2/1	apeq-dc-2000 (not equipped)	up	provisioned	- -	2	N/A

```
=====
```

Sample Output: show peq <peq-slot> detail

```
*A:Dut-A# show peq 1 detail
=====
PEQ 1
=====
```

Slot	Provisioned Type Equipped Type (if diff)	Admin State	Operational State	Input A B	Zone	Input Mode
1	apeq-dc-2000	down	down	Y N	1	N/A

```
-----
Hardware Data
Part number           : 3HE07114AARA01
CLEI code             : IPUPAJHUAA
Serial number         : NS1250G0116
Manufacture date      : 12202012
Manufacturing deviations : (Not Specified)
Manufacturing assembly number: 8205320107
Time of last boot     : 2014/01/07 11:01:44
Current alarm state    : alarm active
=====
```

Sample Output: show peq detail

In the following example, the entries for PEQ 2 to PEQ 11 are not shown.

```
*A:Dut-A# show peq detail
```

```

=====
PEQ 1
=====
Slot   Provisioned Type           Admin Operational   Input   Zone   Input
      Equipped Type (if diff)    State State         A B      Mode
-----
1/1    apeq-dc-2000                down  down             Y N      1    N/A

Hardware Data
  Part number           : 3HE07114AARA01
  CLEI code             : IPUPAJHUAA
  Serial number         : NS1250G0116
  Manufacture date      : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number: 8205320107
  Time of last boot     : 2014/01/07 11:01:44
  Current alarm state    : alarm active

=====

...

=====
PEQ 12
=====
Slot   Provisioned Type           Admin Operational   Input   Zone   Input
      Equipped Type (if diff)    State State         A B      Mode
-----
12     apeq-dc-2000                up    up             Y N      1    N/A

Hardware Data
  Part number           : 3HE07114AARA01
  CLEI code             : IPUPAJHUAA
  Serial number         : NS1250G0115
  Manufacture date      : 12202012
  Manufacturing deviations : (Not Specified)
  Manufacturing assembly number: 8205320107
  Time of last boot     : 2014/01/07 11:01:44
  Current alarm state    : alarm active

=====

```

Output Fields: show peq**Table 69** **Output Fields: show peq**

Label	Description
Slot	The number of the slot in which the APEQ is installed.
Provisioned Type Equipped Type (if different)	The APEQ type provisioned.
Admin State	The administrative state.
Operational State	The operational state.

Table 69 **Output Fields: show peq (Continued)**

Label	Description (Continued)
Input	Specifies the input battery feed: A, or B.
Zone	Specifies the chassis power zone.
Hardware Data:	
Part number	The APEQ part number.
CLEI code	The APEQ CLEI code.
Serial number	The APEQ serial number.
Manufacture date	The date the APEQ was manufactured.
Manufacturing deviations	Specifies any manufacturing deviations.
Manufacturing assembly number	The APEQ assembly number.
Administrative state	Specifies the administrative state of the APEQ.
Operational state	Specifies the operational state of the APEQ.
Time of last boot	Indicates the time stamp of the last system restart.
Current alarm state	Indicates the current alarm state.

megapools

Syntax **megapools** *slot-number* **fp** *fp-number* **wred** [**detail**] **queue-group** *queue-group-name*
 [**instance** *instance-id*]
megapools *slot-number* **fp** *fp-number*
megapools *slot-number* **fp** *fp-number* **wred** [**detail**] [**service-id** *service-id*]

Context show

Description This command displays megapool information. A megapool is a mechanism the forwarding plane uses to allow oversubscription of buffer pools. Every buffer pool is created in the context of a megapool.

By default, all buffer pools are associated with a single megapool and the pools are not oversubscribed. When WRED queue support is enabled on the FP, three megapools are used.

- The original megapool services the default and named pools.
- The second megapool services the system internal use pools.
- The third megapool is used by the buffer pools used by the WRED queues.

The traffic manager buffers are allocated to the three megapools without oversubscription. The WRED queue pools are allowed to oversubscribe their megapool, but the megapool protects the pools associated with the other megapools from buffer starvation that could be caused by that oversubscription.

Parameters	<i>slot-number</i> — Displays information for the specified card slot.
	Values 1 to 10
	<i>fp-number</i> — Displays information for the specified FP.
	Values 1 to 8
	wred — Displays WRED queue pool information.
	detail — Displays detailed information.
	<i>queue-group-name</i> — Displays information for the specified port queue group name, up to 32 characters.
	<i>instance-id</i> — Specifies the identification of a specific instance of the queue group.
	Values 1 to 65535
	<i>service-id</i> — Specifies the service ID.
	Values 1 to 2148278381
	<i>service-name</i> — Specifies a service name up to 64 characters.

sfm

Syntax	sfm [<i>sfm-name</i>] [detail] sfm [<i>sfm-name</i>] icport [down] [degraded]
Context	show
Description	This command displays SFM status information.
Parameters	<i>sfm-name</i> — Specifies the SFM identifier.
	Values 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16
	icport — Displays interconnect port information.
	down — Displays interconnect ports that are not operational.
	degraded — Displays interconnect ports that are associated with fabric degradation.
	detail — Displays detailed information about the SFM.
Output	See the following sections for output samples: • Sample Output: show sfm • Sample Output: show sfm <sfm-id> detail • Sample Output: show sfm icport

- [Output Fields: show sfm](#)

Sample Output: show sfm

```
A:7950 XRS-20# show sfm
=====
SFM Summary
=====
```

Slot	Provisioned Type Equipped Type (if different)	Admin State	Operational State	Comments
1	sfm-x20	up	up	
2	sfm-x20	up	up	
3	sfm-x20	up	up	
4	sfm-x20	up	up	
5	sfm-x20	up	up	
6	(not provisioned) sfm-x20	up	unprovisioned	
7	(not provisioned) sfm-x20	up	unprovisioned	
8	(not provisioned) sfm-x20	up	unprovisioned	

```
=====
```

Sample Output: show sfm <sfm-id> detail

```
A:7950 XRS-20# show sfm 2 detail
=====
Fabric 2
=====
```

Slot	Provisioned Type Equipped Type (if different)	Admin State	Operational State	Comments
2	(not provisioned) sfm-x20	up	unprovisioned	

```
=====
```

Hardware Data

```

Part number           : xx
CLEI code             : xx
Serial number         : xx
Manufacture date      : xx
Manufacturing string   : xx
Manufacturing deviations : xx
Manufacturing assembly number : xx
Administrative state   : up
Operational state     : unprovisioned
Time of last boot     : N/A
Current alarm state    : alarm cleared
=====
```

Inter Chassis SFM Interconnect

```

SFM Interconnect Port 1
oper state : no-link
Misconnect Info : Fabric 3 IcPort 14
SFF Status : not-equipped
fabric degrade state : none
=====
```

Sample Output: show sfm icport

```
*A:myNode# show sfm icport
=====
SFM Interconnect Port Summary
=====
SFM SFM IcPort IcPort Module Degrade Miscon. Info
Oper State Num Oper State Inserted State SFM IcPort
-----
1 unprovisioned 1 up yes none
1 unprovisioned 2 invalid-conne* no degraded 3 14
2 up 2 indeterminate no none
2 up 3 up no degraded
2 up 5 no-link no none
2 up 14 indeterminate yes degraded
=====
* indicates that the corresponding row element may have been truncated.
=====

*A:myNode# show sfm icport down
=====
SFM Interconnect Port Summary
=====
SFM SFM IcPort IcPort Module Degrade Miscon. Info
Oper State Num Oper State Inserted State SFM IcPort
-----
1 unprovisioned 2 invalid-conne* no degraded 3 14
2 up 2 indeterminate no none
2 up 5 no-link no none
2 up 14 indeterminate yes degraded
=====
* indicates that the corresponding row element may have been truncated.
=====

*A:myNode# show sfm icport degraded
=====
SFM Interconnect Port Summary
=====
SFM SFM IcPort IcPort Module Degrade Miscon. Info
Oper State Num Oper State Inserted State SFM IcPort
-----
2 up 3 up no degraded
2 up 14 indeterminate yes degraded
=====
```

Output Fields: show sfm

Table 70 describes output fields for the **show sfm** command.

Table 70 Output Fields: show sfm

Label	Description
Slot	The number of the slot in which the SFM is installed.
Provisioned Type Equipped Type (if different)	The SFM type provisioned.

Table 70 Output Fields: show sfm (Continued)

Label	Description (Continued)
Admin State	The administrative state.
Operational State	The operational state.
Hardware Data	
Part number	The SFM part number.
CLEI code	The SFM CLEI code.
Serial number	The SFM serial number.
Manufacture date	The date the SFM was manufactured.
Manufacturing deviations	Specifies any manufacturing deviations.
Manufacturing assembly number	The SFM assembly number.
Administrative state	Specifies the administrative state of the SFM.
Operational state	Specifies the operational state of the SFM.
Time of last boot	Indicates the time stamp of the last system restart.
Current alarm state	Indicates the current alarm state.
Inter Chassis SFM Interconnect	
SFM Interconnect Port	Port number of the interconnect port.
oper state	Up — The SFM is administratively up. Down — The SFM is administratively down.
Misconnect Info	Only displayed if the oper state is invalid-connection .
SFF Status	The SFF status
fabric degrade state	Indicates state.

2.21.2.3 APS Show Commands

aps

Syntax **aps** [*aps-id*] [*detail*]

Context show

Description This command displays Automated Protection Switching (APS) information.

Parameters *aps-id* — Displays information for the specified APS group ID.

Values *aps-group-id*
aps: keyword
group-id: 1 to 128

detail — Displays detailed APS information.

Output See the following sections for output samples:

- [Sample Output: show aps \(Working Multi-chassis APS Node\)](#)
- [Sample Output: show aps \(Protect MC-APS Node\)](#)
- [Output Fields: show aps](#)

Sample Output: show aps (Working Multi-chassis APS Node)

```
*A:Dut-A# show aps aps-1
=====
APS Group Info
=====
Interface Admin Oper MC-Ctl Work Prot Active Tx/Rx
          State State State Circuit Circuit Circuit K1 Byte
-----
aps-1 Up Up N/A 1/5/1 1/9/5 1/5/1 PC-Tx: No-Req
=====
*A:Dut-A#

*A:Dut-A# show aps aps-1 detail
=====
APS Group: aps-1
=====
Description      : APS Group
Group Id         : 1
Admin Status     : Up
Working Circuit   : 1/5/1
Switching-mode    : Uni-lplus1
Revertive-mode    : Non-revertive
Rx K1/K2 byte     : 0x00/0x00 (No-Req on Protect)
Tx K1/K2 byte     : 0x00/0x00 (No-Req on Protect)
Current APS Status : OK
Multi-Chassis APS : No
Neighbor         : 0.0.0.0
Control link state : N/A
Advertise Interval : 1000 msec
APS SF Hold Time   : 6000 msec
Mode mismatch Cnt  : 0
PSB failure Cnt    : 0
Hold Time         : 3000 msec
APS SD Hold Time   : 9000 msec
Channel mismatch Cnt : 0
FEPL failure Cnt   : 0
-----
APS Working Circuit - 1/5/1
-----
Admin Status      : Up
Current APS Status : OK
Last Switchover    : None
Signal Degrade Cnt : 1
Oper Status       : Up
No. of Switchovers : 0
Switchover seconds : 0
Signal Failure Cnt : 1
```



```

Last Switch Cmd      : No Cmd                      Last Exercise Result : Unknown
Tx L-AIS             : None
-----
APS Protection Circuit - 1/9/5
-----
Admin Status         : Up                          Oper Status          : Up
Current APS Status   : OK                          No. of Switchovers    : 0
Last Switchover      : None                        Switchover seconds    : 0
Signal Degrade Cnt   : 1                          Signal Failure Cnt     : 1
Last Switch Cmd      : No Cmd                      Last Exercise Result   : Unknown
Tx L-AIS             : None
=====
*A:Dut-A#

```

Sample Output: show aps (Protect MC-APS Node)

```

B:Dut-E# show aps
=====
APS Group Info
=====
Interface Admin  Oper    MC-Ctl  Work    Prot    Active    Tx/Rx
                  State    State    State    Circuit  Circuit    Circuit K1 Byte
-----
aps-20    Up    Up    N/A    3/1/1    3/1/2    3/1/1    PC-Tx: No-Req
=====
B:Dut-E#

```

```

B:Dut-E# show aps aps-30 detail
=====
APS Group: aps-30
=====
Description          : APS Group
Group Id              : 30                      Active Circuit        : N/A
Admin Status         : Up                      Oper Status          : Up
Working Circuit       : N/A                    Protection Circuit     : 2/2/2
Switching-mode       : Bi-directional          Switching-arch        : 1+1
Revertive-mode       : Non-revertive            Revert-time (min)    :
Rx K1/K2 byte        : 0x00/0x05 (No-Req on Protect)
Tx K1/K2 byte        : 0x00/0x05 (No-Req on Protect)
Current APS Status   : OK
Multi-Chassis APS    : Yes
Neighbor             : 10.1.1.1
Control link state   : Up
Advertise Interval   : 1000 msec                Hold time             : 3000 msec
Mode mismatch Cnt    : 0                      Channel mismatch Cnt  : 0
PSB failure Cnt      : 0                      FEPL failure Cnt     : 1
-----
APS Working Circuit - Neighbor
-----
Admin Status         : N/A                      Oper Status          : N/A
Current APS Status   : OK                      No. of Switchovers    : 0
Last Switchover      : None                    Switchover seconds    : 0
Signal Degrade Cnt   : 0                      Signal Failure Cnt     : 0
Last Switch Cmd      : No Cmd                  Last Exercise Result   : Unknown
Tx L-AIS             : None

```

```

-----
APS Protection Circuit - 2/2/2
-----
Admin Status      : Up                Oper Status      : Up
Current APS Status : OK                No. of Switchovers : 0
Last Switchover    : None              Switchover seconds : 0
Signal Degrade Cnt : 0                Signal Failure Cnt  : 0
Last Switch Cmd    : No Cmd            Last Exercise Result : Unknown
Tx L-AIS           : None
=====
B:Dut-E#

```

Output Fields: show aps

[Table 71](#) describes output fields for the **show aps** command.

Table 71 **Output Fields: show aps**

Label	Description
Interface	Specifies the APS interface name (the APS group port).
Admin State	Up — APS is administratively up. Down — APS port is administratively down.
Oper State	Up — APS port is operationally up. Down — APS is operationally down.
MC-CTL State	Specifies the multi-chassis state.
Work Circuit	Specifies the working circuit ID.
Prot Circuit	Specifies the physical port that acts as the protection circuit for this APS group.
Active Circuit	Specifies the active circuit.
Tx/Rx K1 Byte	Displays the value of the SONET/SDH K1 byte received or transmitted on the protection circuit.
Group Id	Displays the APS group name.
Protection Circuit	Displays the physical port that will act as the protection circuit for this APS group.
Switching-mode	Displays the switching mode of the APS group.
Switching-arch	The architecture of the APS group.
Revertive-mode	Displays the revertive mode of the APS group. non-revertive — Traffic remains on the protection line until another switch request is received. revertive — When the condition that caused a switch to the protection line has been cleared the signal is switched back to the working line.

Table 71 **Output Fields: show aps (Continued)**

Label	Description (Continued)
Revert-time	Displays the configured time, in minutes, to wait after the working circuit has become functional again, before making the working circuit active again. If the revertive mode is non-revertive, then this field will be empty.
Rx K1/K2 byte	Displays the value of the SONET/SDH K1/K2 byte received on the interface.
Tx K1/K2 byte	Displays the value of the SONET/SDH K1/K2 byte transmitted on the interface.
Current APS Status	Displays the current APS status.
Mode Mismatch Cnt	Indicates the number of times a conflict occurs between the current local mode and the received K2 mode information.
Channel mismatch Cnt	Indicates the number of mismatches between the transmitted K1 channel and the received K2 channel has been detected.
PSB failure Cnt	Displays a count of Protection Switch Byte (PSB) failure conditions. This condition occurs when either an inconsistent APS byte or an invalid code is detected.
FEPL failure Cnt	Displays a count of far-end protection-line (FEPL) failure conditions. This condition is declared based on receiving SF on the protection line in the K1 byte.
No. of Switchovers	Displays the number of times a switchover has occurred.
Last Switchover	Displays the time stamp of the last switchover.
Switchover seconds	Displays the cumulative Protection Switching Duration (PSD) time in seconds. For a working channel, this is the cumulative number of seconds that service was carried on the protection line. For the protection line, this is the cumulative number of seconds that the protection line has been used to carry any working channel traffic. This information is only valid if revertive switching is enabled.
Signal Degrade Cnt	Displays the number of times the signal was degraded.
Signal Failure Cnt	Displays the number of times the signal failed.
Last Switch Cmd	Reports the last switch command that was performed on a circuit.
Last Exercise Result	The result of the last exercise request on a circuit.
Neighbor address	Displays the neighbor IP address.
Advertise Interval	Displays the advertise interval.
Hold time	Displays the hold time.

2.21.2.4 Port Show Commands

port

```
Syntax
port [port-id] [statistics [egress-aggregate]] [detail]
port port-id associations
port [port-id] description
port port-id dotx1 [detail]
port aps
port [port-id] cem
port port-id atm [detail]
port port-id atm connections
port port-id atm cp [cp] [detail]
port port-id atm ilmi [detail]
port port-id atm interface-connection [detail]
port port-id atm pvc [vpi[/vci]] [detail]
port port-id atm pvp [vpi] [detail]
port port-id atm pvt [vpi1.vpi2] [detail]
port port-id cisco-hlhc
port port-id ethernet [association | detail]
port port-id ethernet [association | detail] efm-oam [event-logs [{failure | degraded}]
    [{active | cleared}]]
port port-id ethernet [association | detail] lldp [nearest-bridge | nearest-non-tpmr |
    nearest-customer] [remote-info] [detail]
port port-id frame-relay [detail]
port port-id ima-link
port port-id macsec sub-port sub-port-id [detail] [statistics]
port port-id mlfr-link
port port-id monitor-threshold
port port-id optical
port port-id otu
port port-id ppp [detail]
port port-id queue-group [ingress | egress] [queue-group-name] [access | network]
    [instance instance-id]
port port-id queue-group [ingress | egress] [queue-group-name] [access | network]
    associations [instance instance-id]
port port-id queue-group [ingress | egress] queue-group-name [access | network]
    [instance instance-id] queue-depth [queue queue-id]
port port-id queue-group [ingress | egress] [queue-group-name] [access | network]
    statistics [instance instance-id]
port port-id queue-group summary
port port-id vport vport-name monitor-threshold
```

port *port-id* **wavekey-table** {50g | 100g}
port *port-id* **wavelength-table**
port *port-id port-id* **wavetracker**

Context show

Description This command displays port or channel information.

If the *port-id* parameter only specifies a portion of a port identifier then a summary of all ports that start with that portion is displayed. For example, specifying a slot number and mda number will present a summary of all ports on that mda. If no *port-id* is provided then a summary of all ports in the system is displayed.

If the **detail** keyword is specified without any *port-id* then the detailed output of every port is displayed. This is useful for gathering full reporting or as an input into the "|" match" post-filtering command to display only certain fields for all ports (customized summaries).

If the *port-identifier* specifies a unique port, then information about that port is displayed. The specific information displayed depends on the type of port. Additional information can be displayed if the **detail** is included. In addition, information subsets can be presented by using various keywords. For example, the **otu** keyword displays only the OTU Interface information for the port.

Parameters *port-id* — Specifies the physical port ID in the form *slot/mda/port*.

Values

<i>port-id</i>	slot[/mda[/port]] or slot/mda/port [.channel]	
aps-id	aps-group-id[.channel]	
	aps	keyword
	group-id	1 to 64
ccag-id	slot/mda/path-id[cc-type]	
	path-id	a,b
	cc-type	.sap-net, .net-sap
eth-sat-id	esat-id [/slot/[u]port]	
	esat	keyword
	id	1 to 20
	u	keyword for up-link port
pxc-id	pxc-id.sub-port	
	pxc	keyword
	id	1 to 64
	sub-port	a, b

Table 72 Port ID Values

Product	Slot	MDA	CMA	Port	Values
7750 SR-12	1 to 10	1, 2		1 to 60 (depending on the MDA type)	—
7750 SR-c12	1	1, 3, 5, 7, 9, 11	1 to 12		—
7750 SR-c4	1	1, 3	1 to 4		—
7750 SR-7	1 to 5	1, 2	—		—
7950 XRS	1 to 20	—	—		—
7450 ESS-7	—		—		1 to 4
7450 ESS-12	—		—		1 to 0
Channelized MDAs					
CHOC12-SFP	—				slot/mda/port. [1 to 4]. [1 to 3]. [1 to 28]. [..24] For example, 7/2/1.1.1.28.24
CHOC3-SFP	—				slot/mda/port. [1 to 3]. [1 to 28]. [..24] For example, 7/2/1.1.1.28.24
DS3	—				slot/mda/port. [1 to 28] . [..24] For example, 7/1/1.1.1

statistics — Displays port statistics.

egress-aggregate — Displays aggregate egress statistics.

detail — Provides detailed information.

associations — Displays a list of current router interfaces to which the port is associated.

description — Displays port description strings.

dot1x — Displays information about 802.1x status and statistics.

aps — Displays ports on APS groups.

cem — Displays CEM encap ports and channels.

atm — Displays ATM information.

connections — Displays ATM connection information.

cp — Displays ATM port connection profile information.

ilmi — Displays ATM ILMI information.

interface-connection — Displays ATM connection information.

pvc — Displays ATM port PVC information.

pvp — Displays ATM port PVP information.

pvt — Displays ATM port PVT information.

vci — Specifies the ATM network virtual channel identifier (VCI).

Values 1, 2, 5 to 65535

vpi — Specifies the ATM network virtual path identifier (VPI).

Values 0 to 4095 (NNI)

0 to 255 (UNI)

vpi1.vpi2 — Specifies the ATM network virtual path identifiers (VPIs).

Values 0 to 4095.0 to 4095 (NNI)

0 to 255.0 to 255 (UNI)

cisco-hdlc — Displays Cisco HDLC port information.

ethernet — Displays Ethernet port information.

failure — Keyword to display failure severity events.

degraded — Keyword to display degrade severity events.

active — Keyword to display only active events.

cleared — Keyword to display remote information on the bridge MAC.

frame-relay — Displays Frame Relay information.

nearest bridge — Displays nearest bridge information.

nearest-non-tpmr — Displays nearest Two-Port MAC Relay (TPMR) information.

nearest-customer — Displays nearest customer information.

remote-info — Displays remote information on the bridge MAC.

ima-link — Displays the link-based IMA information for the port.

macsec — Displays the MACsec information for the port.

mlfr-link — Displays link-based MLFR information for the port.

monitor-threshold — Displays the exceed-count for the port-scheduler under Vport (if specified) or for a physical port.

optical — Displays optical information.

otu — Displays optical transport unit (OTU) information.

ppp — Displays PPP protocol information for the port.

queue-group — Displays the queue group information.

queue-depth — Displays the queue depth information.

instance-id — Specifies the identification of a specific instance of the queue group.

Values 1 to 65535

vport-name — Specifies the name of the Vport up to 32 characters in length.

associations — Displays a list of ports to which the Vport is assigned.

wavekey-table — Displays a table of acceptable wave-keys for the different DWDM channels.

50g | 100g — Displays specified DWDM wave key table information.

wavelength-table — Displays DWDM wave length table information.

wavetracker — Displays Wave Tracker information.

Output See the following sections for output samples:

- [Sample Output: show port <port-id> \(Summary Table of Ports\)](#)
- [Output Fields: show port <port-id>](#)
- [Sample Output: show port <port-id> \(Showing Ethernet Interface\)](#)
- [Sample Output: show port <port-id> detail \(Ethernet Interface Port\)](#)
- [Output Fields: show port <port-id> detail](#)
- [Sample Output: show port <port id> optical detail](#)
- [Output Fields: show port <port-id> optical detail](#)
- [Sample Output: show port <port-id> vport <vport-name>](#)
- [Sample Output: show port <port-id> detail \(Excerpt Showing Ethernet Statistics\)](#)
- [Output Fields: show port <port-id> detail \(Excerpt Showing Ethernet Statistics\)](#)
- [Sample Output: show port <port-id> detail \(Excerpt Showing Ethernet-like Medium Statistics\)](#)
- [Output Fields: show port <port-id> detail \(Excerpt Showing Ethernet-like Medium Statistics Output\)](#)
- [Sample Output: show port <port-id> \(Showing Channelized Ports\)](#)
- [Output Fields: show port <port-id> \(Showing Channelized Ports\)](#)
- [Sample Output: show port <port-id> associations](#)
- [Output Fields: show port <port-id> associations](#)
- [Sample Output: show port <port-id> frame-relay](#)
- [Output Fields: show port <port-id> frame-relay](#)
- [Sample Output: show port <port-id> otu detail](#)
- [Output Fields: show port <port-id> otu detail](#)
- [Sample Output: show port <port-id> ppp](#)
- [Output Fields: show port <port-id> ppp](#)
- [Sample Output: show port <port-id> atm](#)
- [Sample Output: show port <port-id> atm cp](#)
- [Output Fields: show port <port-id> atm](#)
- [Sample Output: show port <port-id> atm pvc detail](#)
- [Output Fields: show port <port-id> atm pvc detail](#)
- [Sample Output: show port <port-id> atm pvt detail](#)

- Output Fields: show port <port-id> atm pvt detail
- Sample Output: show port <port-id> [statistics [egress-aggregate]] [detail]
- Sample Output: show port statistics
- Sample Output: show port <port-id> dot1x
- Output Fields: show port <port-id> dot1x
- Sample Output: show port esat-n/n/n
- Sample Output: show port ethernet
- Sample Output: show port ethernet efm-oam
- Sample Output: show port ethernet efm-oam event-logs
- Sample Output: show port <port-id> ethernet lldp
- Sample Output: show port <port-id> ethernet lldp remote-info
- Sample Output: show port <port-id> ethernet lldp remote-info detail
- Sample Output: show port <port-id> ethernet lldp detail
- Sample Output: show port <port-id> macsec sub-port <sub-port-id>
- Output Fields: show port <port-id> macsec subport-id <sub-port-id>
- Sample Output: show port <port-id> macsec sub-port <sub-port-id> detail
- Sample Output: show port <port-id> macsec sub-port <sub-port-id> statistics
- Output Fields: show port <port-id> macsec sub-port <sub-port-id> statistics

Sample Output: show port <port-id> (Summary Table of Ports)

```
*A:ALU-1# show port 1/1
=====
Ports on Slot 1
=====
Port      Admin Link Port   Cfg  Oper LAG/  Port Port Port   SFP/XFP/
Id        State  State  MTU  MTU  Bndl Mode Encp Type  MDIMDX
-----
1/1/1     Down  No   Down  1518 1518   1 accs dotq gige
1/1/2     Down  No   Down  1578 1578   - netw null gige
1/1/3     Down  No   Down  1578 1578   - netw null gige
1/1/4     Up    No   Down  1514 1514   - accs null gige
1/1/5     Up    No   Down  1578 1578   - netw null gige
=====
*A:ALU-1#
```

Output Fields: show port <port-id>

[Table 73](#) describes the output fields for the **show port <port-id>** command.

Table 73 Output Fields: show port <port-id>

Label	Description
Port ID	The port ID configured or displayed.

Table 73 **Output Fields: show port <port-id> (Continued)**

Label	Description (Continued)
Admin State	Up — The administrative state is up. Down — The administrative state is down.
Phy Link	Yes — A physical link is present. No — A physical link is not present.
Port State	Up — The port is physically present and has physical link present. Down — The port is physically present but does not have a link. Note that this state may also be considered as Link Down. Ghost — A port that is not physically present. None — The port is in its initial creation state or about to be deleted. Link Up — A port that is physically present and has physical link present. Note that when Link Up appears at the lowest level of a SONET/SDH path or a TDM tributary, it means the physical connection is active but the port is waiting on some other state before data traffic can flow. It is a waiting state and indicates that data traffic will not flow until it transitions to the Up state.
Cfg MTU	The configured MTU.
Oper MTU	The negotiated size of the largest packet which can be sent on the port SONET/SDH, channel, specified in octets. For channels that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the channel.
LAG ID	The LAG or multi-link trunk (MLT) that the port is assigned to.
Port Mode	Network — The port is configured for transport network use. Access — The port is configured for service access. Hybrid — The port is configured for both access and network use.
Port Encap	Null — Ingress frames will not use tags or labels to delineate a service. dot1q — Ingress frames carry 802.1Q tags where each tag signifies a different service.
Port Type	The type of port or optics installed.
SFP/MDI MDX	GIGE — Indicates the GigE SFP type. FASTE — Indicates the FastE SFP type. GIGX — Indicates the GigX SFP type. MDI — Indicates that the Ethernet interface is of type MDI (Media Dependent Interface). MDX — Indicates that the Ethernet interface is of type MDX (Media Dependent Interface with crossovers).

Sample Output: show port <port-id> (Showing Ethernet Interface)

```
*A:ALU-211# show port 1/1/2
=====
Ethernet Interface
=====
Description      : 10/100 Ethernet TX
Interface        : 1/1/2                      Oper Speed      : 100 mbps
Link-level       : Ethernet                  Config Speed    : 100 mbps
Admin State      : up                       Oper Duplex     : full
Oper State       : up - Active in LAG 10     Config Duplex   : full
Physical Link    : Yes                      MTU             : 1514
Single Fiber Mode : No
IfIndex          : 35717120                 Hold time up    : 0 seconds
Last State Change : 12/16/2008 19:31:40      Hold time down  : 0 seconds
Last Cleared Time : 12/16/2008 19:31:48
.....
=====
*A:ALU-211#
```

```
*A:ALU-1# show port 5/1/4
=====
Ethernet Interface
=====
Description      : 10-Gig Ethernet
Interface        : 5/1/4                      Oper Speed      : 10 Gbps
Link-level       : Ethernet                  Config Speed    : N/A
Admin State      : up                       Oper Duplex     : full
Oper State       : down                    Config Duplex   : N/A
Physical Link    : No                      MTU             : 9212
Single Fiber Mode : No                     Min Frame Length : 64 Bytes
IfIndex          : 170000384                Hold time up    : 0 seconds
Last State Change : 04/28/2017 13:09:15      Hold time down  : 0 seconds
Last Cleared Time : N/A                     DDM Events     : Enabled
Phys State Chng Cnt: 10                     RS-FEC Mode    : None

Configured Mode   : network                  Encap Type      : null
Dot1Q Ethertype   : 0x8100                  QinQ Ethertype  : 0x8100
PBB Ethertype     : 0x88e7
Ing. Pool % Rate  : 100                      Egr. Pool % Rate : 100
Ing. Pool Policy  : n/a
Egr. Pool Policy  : n/a
Net. Egr. Queue Pol : default
Egr. Sched. Pol   : n/a
HS Scheduler Plcy : default
HS Port Pool Plcy : default
Monitor Port Sched : Disabled
Monitor Agg Q Stats: Disabled
Auto-negotiate    : N/A                     MDI/MDX         : N/A
Oper Phy-tx-clock : not-applicable
Accounting Policy : None                     Collect-stats   : Disabled
Acct Plcy Eth Phys : None                    Collect Eth Phys : Disabled
Egress Rate       : Default                  Ingress Rate    : Default
Load-balance-algo : Default                  LACP Tunnel     : Disabled
Access Bandwidth  : Not-Applicable           Booking Factor   : 100
Access Available BW: 0
Access Booked BW  : 0
Sflow             : Disabled
```

```

Suppress Threshold : 1500
Max Penalties     : 4000
Half Life         : 20 seconds

Down-when-looped  : Disabled
Loop Detected     : False
Use Broadcast Addr : False

Sync. Status Msg. : Disabled
Tx DUS/DNU        : Disabled
SSM Code Type     : sdh

Down On Int. Error : Disabled

CRC Mon SD Thresh : Disabled
CRC Mon SF Thresh : Disabled

Sym Mon SD Thresh : Disabled
Sym Mon SF Thresh : Disabled

EFM OAM           : Disabled
Ignr EFM OAM State : False

Configured Address : 10:e8:78:4d:24:f0
Hardware Address   : 10:e8:78:4d:24:f0
Cfg Alarm          : remote local

Reuse Threshold    : 1000
Max Suppress Time : 40 seconds

Keep-alive         : 10
Retry              : 120

Rx Quality Level   : N/A
Tx Quality Level   : N/A

DOIE Tx Disable    : Disabled

CRC Mon Window     : 10 seconds

Sym Mon Window     : 10 seconds
Tot Sym Mon Errs   : 0

EFM OAM Link Mon   : Disabled

```

Transceiver Data

```

Transceiver Status : operational
Transceiver Type    : SFP
Model Number        : 3HE04823AAAA01 ALA IPU3ANKEAA
TX Laser Wavelength: 1310 nm
Connector Code      : LC
Manufacture date    : 2016/07/23
Serial Number       : AW40PC8
Part Number         : FTLX1471D3BCL-A5
Optical Compliance  : 10GBASE-LR
Link Length support: 10km for SMF

Diag Capable        : yes
Vendor OUI          : 00:90:65
Media               : Ethernet

```

```

=====
Transceiver Digital Diagnostic Monitoring (DDM), Internally Calibrated
=====

```

	Value	High Alarm	High Warn	Low Warn	Low Alarm
Temperature (C)	+38.1	+78.0	+73.0	-8.0	-13.0
Supply Voltage (V)	3.31	3.70	3.60	3.00	2.90
Tx Bias Current (mA)	40.4	85.0	80.0	20.0	15.0
Tx Output Power (dBm)	-1.30	2.00	1.00	-7.00	-8.00
Rx Optical Power (avg dBm)	-1.38	2.50	2.00	-18.01	-20.00

```

=====
OTU Interface
=====

```

```

OTU Status          : Disabled
=====

```

```

=====
Traffic Statistics
=====
                                     Input          Output
-----
Octets                             0              0
Packets                            0              0
Errors                             0              0
Utilization (300 seconds)          0.00%          0.00%

=====
Port Statistics
=====
                                     Input          Output
-----
Unicast Packets                     0              0
Multicast Packets                    0              0
Broadcast Packets                    0              0
Discards                            0              0
Unknown Proto Discards               0

=====
Ethernet-like Medium Statistics
=====
Alignment Errors :                   0  Sngl Collisions :                   0
FCS Errors       :                   0  Mult Collisions :                   0
SQE Test Errors  :                   0  Late Collisions :                   0
CSE              :                   0  Excess Collisns :                   0
Too long Frames  :                   0  Int MAC Tx Errs :                   0
Symbol Errors    :                   0  Int MAC Rx Errs :                   0
In Pause Frames  :                   0  Out Pause Frames :                   0
=====

```

Sample Output: show port <port-id> detail (Ethernet Interface Port)

```

A:ALU-1# show port 5/1/4 detail
=====
Ethernet Interface
=====
Description      : 10-Gig Ethernet
Interface        : 5/1/4
Link-level       : Ethernet
Admin State      : up
Oper State       : down
Physical Link     : No
Single Fiber Mode : No
IfIndex          : 170000384
Last State Change : 04/28/2017 13:09:15
Last Cleared Time  : N/A
Phys State Chng Cnt: 10

Oper Speed       : 10 Gbps
Config Speed     : N/A
Oper Duplex      : full
Config Duplex    : N/A
MTU              : 9212
Min Frame Length : 64 Bytes
Hold time up     : 0 seconds
Hold time down   : 0 seconds
DDM Events       : Enabled
RS-FEC Mode      : None

Configured Mode   : network
Dot1Q Ethertype   : 0x8100
PBB Ethertype     : 0x88e7
Ing. Pool % Rate  : 100
Ing. Pool Policy  : n/a
Egr. Pool Policy  : n/a
Net. Egr. Queue Pol: default

Encap Type       : null
QinQ Ethertype    : 0x8100
Egr. Pool % Rate  : 100

```

```

Egr. Sched. Pol      : n/a
HS Scheduler Plcy    : default
HS Port Pool Plcy    : default
Monitor Port Sched   : Disabled
Monitor Agg Q Stats  : Disabled
Auto-negotiate       : N/A
Oper Phy-tx-clock    : not-applicable
Accounting Policy     : None
Acct Plcy Eth Phys   : None
Egress Rate          : Default
Load-balance-algo    : Default
Access Bandwidth     : Not-Applicable
Access Available BW   : 0
Access Booked BW     : 0
Sflow                : Disabled

MDI/MDX              : N/A

Collect-stats        : Disabled
Collect Eth Phys     : Disabled
Ingress Rate         : Default
LACP Tunnel          : Disabled
Booking Factor       : 100

Dampening State      : Active
Suppress Threshold   : 1500
Max Penalties        : 4000
Half Life            : 20 seconds

Current Penalties    : 2297
Reuse Threshold      : 1000
Max Suppress Time    : 40 seconds

Down-when-looped     : Disabled
Loop Detected        : False
Use Broadcast Addr    : False

Keep-alive           : 10
Retry                : 120

Sync. Status Msg.    : Disabled
Tx DUS/DNU           : Disabled
SSM Code Type        : sdh

Rx Quality Level     : N/A
Tx Quality Level     : N/A

Down On Int. Error   : Disabled
DOIE Tx Disable      : Disabled

CRC Mon SD Thresh    : Disabled
CRC Mon SF Thresh    : Disabled

CRC Mon Window       : 10 seconds

Sym Mon SD Thresh    : Disabled
Sym Mon SF Thresh    : Disabled

Sym Mon Window       : 10 seconds
Tot Sym Mon Errs     : 0

EFM OAM              : Disabled
Ignr EFM OAM State   : False

EFM OAM Link Mon     : Disabled

Configured Address   : 10:e8:78:4d:24:f0
Hardware Address     : 10:e8:78:4d:24:f0
Cfg Alarm            : remote local

Transceiver Data

Transceiver Status   : operational
Transceiver Type     : SFP
Model Number         : 3HE04823AAAA01 ALA IPU3ANKEAA
TX Laser Wavelength : 1310 nm
Connector Code       : LC
Manufacture date     : 2016/07/23
Serial Number        : AW40PC8
Part Number          : FTLX1471D3BCL-A5
Optical Compliance   : 10GBASE-LR
Link Length support  : 10km for SMF

Diag Capable         : yes
Vendor OUI           : 00:90:65
Media                : Ethernet

=====
Transceiver Digital Diagnostic Monitoring (DDM), Internally Calibrated

```

```
=====
```

	Value	High Alarm	High Warn	Low Warn	Low Alarm
Temperature (C)	+38.1	+78.0	+73.0	-8.0	-13.0
Supply Voltage (V)	3.31	3.70	3.60	3.00	2.90
Tx Bias Current (mA)	40.3	85.0	80.0	20.0	15.0
Tx Output Power (dBm)	-1.29	2.00	1.00	-7.00	-8.00
Rx Optical Power (avg dBm)	-1.38	2.50	2.00	-18.01	-20.00

```
=====
```

```
=====
```

OTU Interface

```
=====
```

OTU Status : Disabled

```
=====
```

```
=====
```

Traffic Statistics

```
=====
```

	Input	Output
Octets	0	0
Packets	0	0
Errors	0	0
Utilization (300 seconds)	0.00%	0.00%

```
=====
```

```
=====
```

Ethernet Statistics

```
=====
```

Broadcast Pckts :	0	Drop Events :	0
Multicast Pckts :	0	CRC/Align Errors :	0
Undersize Pckts :	0	Fragments :	0
Oversize Pckts :	0	Jabbers :	0
Collisions :	0		
Octets :	0		
Packets :	0		
Packets of 64 Octets :	0		
Packets of 65 to 127 Octets :	0		
Packets of 128 to 255 Octets :	0		
Packets of 256 to 511 Octets :	0		
Packets of 512 to 1023 Octets :	0		
Packets of 1024 to 1518 Octets :	0		
Packets of 1519 or more Octets :	0		

```
=====
```

```
=====
```

Port Statistics

```
=====
```

	Input	Output
Unicast Packets	0	0
Multicast Packets	0	0
Broadcast Packets	0	0
Discards	0	0
Unknown Proto Discards	0	

```
=====
```

Ethernet-like Medium Statistics

```

=====
Alignment Errors :                0  Sngl Collisions :                0
FCS Errors       :                0  Mult Collisions :                0
SQE Test Errors  :                0  Late Collisions :                0
CSE              :                0  Excess Collisns :                0
Too long Frames  :                0  Int MAC Tx Errs :                0
Symbol Errors    :                0  Int MAC Rx Errs :                0
In Pause Frames  :                0  Out Pause Frames :                0
=====

```

Per Threshold MDA Discard Statistics

```

=====
                                Packets                Octets
-----
Threshold 0 Dropped :                0                0
Threshold 1 Dropped :                0                0
Threshold 2 Dropped :                0                0
Threshold 3 Dropped :                0                0
Threshold 4 Dropped :                0                0
Threshold 5 Dropped :                0                0
Threshold 6 Dropped :                0                0
Threshold 7 Dropped :                0                0
Threshold 8 Dropped :                0                0
Threshold 9 Dropped :                0                0
Threshold 10 Dropped :               0                0
Threshold 11 Dropped :               0                0
Threshold 12 Dropped :               0                0
Threshold 13 Dropped :               0                0
Threshold 14 Dropped :               0                0
Threshold 15 Dropped :               0                0
=====

```

Ingress Port Forwarding Engine Drop Reason Statistics

```

=====
IPv4 Header Error                0
IPv4 Invalid Address              0
IPv6 Header Error                0
IPv6 Invalid Address              0
IP Route Blackholed               0
ACL Filter Discards               0
Unicast RPF Check Failed          0
BFDD Spoof Check Failed           0
Unicast MAC Destination Address Mismatch 0
Multicast MAC With Unicast Dest IP 0
Unknown MAC Destination Address Discarded in VPLS 0
L2 Service MTU Exceeded           0
Needs ICMP                        0
=====

```

Queue Statistics

```

-----
Ingress Queue 1                Packets                Octets
In Profile forwarded :         0                0

```

```

        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 2             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 3             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 4             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 5             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 6             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 7             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 8             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 9             Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 10            Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 11            Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 12            Packets      Octets
        In Profile forwarded    :      0      0
        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
    Ingress Queue 13            Packets      Octets

```

```

        In Profile forwarded :    0          0
        In Profile dropped   :    0          0
        Out Profile forwarded :    0          0
        Out Profile dropped  :    0          0
    Ingress Queue 14         Packets          Octets
        In Profile forwarded :    0          0
        In Profile dropped   :    0          0
        Out Profile forwarded :    0          0
        Out Profile dropped  :    0          0
    Ingress Queue 15         Packets          Octets
        In Profile forwarded :    0          0
        In Profile dropped   :    0          0
        Out Profile forwarded :    0          0
        Out Profile dropped  :    0          0
    Ingress Queue 16         Packets          Octets
        In Profile forwarded :    0          0
        In Profile dropped   :    0          0
        Out Profile forwarded :    0          0
        Out Profile dropped  :    0          0

    Egress Queue  1         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  2         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  3         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  4         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  5         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  6         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  7         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0
        Out/Exc Prof dropped :    0          0
    Egress Queue  8         Packets          Octets
        In/Inplus Prof fwded :    0          0
        In/Inplus Prof dropped:    0          0
        Out/Exc Prof fwded   :    0          0

```

```
Out/Exc Prof dropped :      0
=====
```

Output Fields: show port <port-id> detail

[Table 74](#) describes the output fields for the **show port <port-id> detail** command.

Table 74 Output Fields: show port <port-id> detail

Label	Description
Description	A text description of the port.
Interface	The port ID displayed in the <i>slot/mda/port</i> format.
Oper Speed	The operating speed of the interface.
Link-level	Ethernet — The port is configured as Ethernet. SONET — The port is configured as SONET-SDH.
Config Speed	The configured speed of the interface.
Admin State	up — The port is administratively up. down — The port is administratively down.
Oper Duplex	full — The link is set to full duplex mode. half — The link is set to half duplex mode.
Oper State	up — The port is operationally up. down — The port is operationally down. Additionally, the <i>lag-id</i> of the LAG it belongs to in addition to the status of the LAG member (active or standby) is specified.
Config Duplex	Full — The link is set to full duplex mode. Half — The link is set to half duplex mode.
Physical Link	Yes — A physical link is present. No — A physical link is not present.
MTU	The size of the largest packet which can be sent/received on the Ethernet physical interface, specified in octets.
Single Fiber Mode	Yes - Single fiber option is configured. No - Single fiber option is not configured.
Min Frame Length	Displays the configured minimum transmitted frame length.
IfIndex	Displays the interface's index number which reflects its initialization sequence.
Hold time up	The link up dampening time in seconds. The port link dampening timer value which reduces the number of link transitions reported to upper layer protocols.

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
Last State Change	Displays the system time moment that the MC-LAG group is up.
Hold time down	The link down dampening time in seconds. The down timer controls the dampening timer for link down transitions.
Last Cleared Time	Displays the system time moment that the peer is up.
DDM Events	Enabled — DDM events are enabled. Disabled — DDM events are disabled.
Phys State Chng Cnt	Increments when a fully qualified (de-bounced) transition occurs at the physical layer of an ethernet port which includes the following transitions of the Port State as shown in the "show port" summary: - from "Down" to either "Link Up" or "Up" - from either "Link Up" or "Up" to "Down" This counter does not increment for changes purely in the link protocol states (e.g. "Link Up" to "Up"). The counter is reset if the container objects for the port are deleted (e.g. MDA deconfigured, or IOM type changes).
RS-FEC Mode	Specifies the RS-FEC mode to support.
Configured Mode	network — The port is configured for transport network use. access — The port is configured for service access.
Encap Type	Null — Ingress frames will not use any tags or labels to delineate a service. dot1q — Ingress frames carry 802.1Q tags where each tag signifies a different service.
Dot1Q Ethertype	Indicates the Ethertype expected when the port's encapsulation type is Dot1Q.
QinQ Ethertype	Indicates the Ethertype expected when the port's encapsulation type is QinQ.
PBB Ethertype	Indicates the Ethertype used for PBB encapsulation.
Net. Egr. Queue Pol	Specifies the network egress queue policy or that the default policy is used.
Egr. Sched. Pol	Specifies the port scheduler policy or that the default policy default is in use.
Monitor Port Sched	Enabled — Congestion monitoring on an Egress Port Schedule (EPS) is enabled. Disabled — Congestion monitoring on an EPS is disabled.
Auto-negotiate	True — The link attempts to automatically negotiate the link speed and duplex parameters. False — The duplex and speed values are used for the link.

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
Collect-stats	Enabled — The collection of accounting and statistical data for the network Ethernet port is enabled. When applying accounting policies the data by default will be collected in the appropriate records and written to the designated billing file. Disabled — Collection is disabled. Statistics are still accumulated by the IOM cards, however, the CPU will not obtain the results and write them to the billing file.
Egress Rate	The maximum amount of egress bandwidth (in kilobits per second) that this Ethernet interface can generate.
Ingress Rate	Indicates the maximum amount of ingress bandwidth (in mb/s) that this Ethernet port can receive with the configured sub-rate using packet-based accounting.
Load-balance-algo	Indicates the load balancing algorithm used on the port.
LACP Tunnel	Indicates whether LACP packet tunneling is enabled or disabled.
Sflow	Enabled — sFlow data collection for the port is enabled. Disabled — sFlow data collection is disabled..
Dampening State	Displays the port dampening state: Disabled, Idle, or Active
Current Penalties	Displays the current accumulated penalties for port dampening.
Suppress Threshold	Displays the threshold at which the port-up state is suppressed until the accumulated penalties drop below the reuse threshold again.
Reuse Threshold	Displays the threshold at which the port-up state is no longer suppressed, after the port has been in a suppressed state and the accumulated penalties decay drops below this threshold.
Max Penalties	Displays the maximum penalty value for port dampening.
Max Suppress Time	Displays the time, in seconds, it can take after the link comes up before the worst case accumulated penalties have decayed to the reuse threshold.
Half Life	Displays the time, in seconds, that must pass before penalties decay to one-half the initial amount.
Down-when-looped	Shows whether the feature is enabled or disabled.
Keep Alive	Displays the time interval, in seconds, between keep-alive PDUs.
Loop Detected	True — Loop detected. False — No loop detected.
Retry	Displays the minimum wait time, in seconds, before re-enabling the port after loop detection.

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
Use Broadcast Addr	True — Indicates that the broadcast address is to be used for the destination MAC address. False — Indicates that the local port MAC address is to be used for the destination MAC address.
Sync. Status Msg	Whether synchronization status messages are enabled or disabled.
Rx Quality Level	Indicates which QL value has been received from the interface.
Tx DUS/DNU	Whether the QL value is forcibly set to QL-DUS/QL-DNU.
Tx Quality Level	Indicates which QL value is being transmitted out of the interface.
SSM Code Type	Indicates the SSM code type in use on the port.
Down On Int. Error	Enabled — Down on internal error feature is enabled. Disabled — Down on internal error feature is disabled.
DOIE Tx Disable	Enabled — Laser is enabled if an internal MAC transmit error is encountered. Disabled — Laser is disabled if an internal MAC transmit error is encountered.
CRC Mon SD Thresh	Specifies the error rate (for CRC errors) at which to declare the Signal Degrade (SD) condition on an Ethernet interface. If the field displays Disabled, no error rate has been specified.
CRC Mon Window	Specifies the size of the sliding window, in seconds, over which the Ethernet frames are sampled to detect signal fail or signal degrade conditions.
CRC Mon SF Thresh	Specifies the CRC error rate at which to declare the Signal Fail (SF) condition on an Ethernet interface. If the field displays Disabled, no error rate has been specified.
Sym Mon SD Thresh	Specifies the error rate (for symbol errors) at which to declare the Signal Degrade (SD) condition on an Ethernet interface. If the field displays Disabled, no error rate has been specified.
Sym Mon Window	Indicates the size of the sliding window, in seconds, over which the errors are measured.
Sym Mon SF Thresh	Specifies the symbol error rate at which to declare the Signal Fail (SF) condition on an Ethernet interface. If the field displays Disabled, no error rate has been specified.
Tot Sym Mon Errs	Displays the total number of symbol errors.
EFM OAM	Enabled — EFM OAM is enabled. Disabled — EFM OAM is disabled.
EFM OAM Link Mon	Enabled — Link monitoring functionality is enabled. Disabled — Link monitoring functionality is disabled.

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
Ignr EFM OAM State	Enabled — Any failure in the protocol state machine does not impact the state of the port. Disabled — The port state is affected by any existing EFM-OAM protocol fault condition
Configured Address	The base chassis Ethernet MAC address.
Hardware Address	The interface's hardware or system assigned MAC address at its protocol sub-layer.
Transceiver Data	See Table 75 .
Transceiver Status	Status of the transceiver.
Transceiver Type	Type of the transceiver.
Model Number	The model number of the transceiver.
TX Laser Wavelength	Indicates the transceiver laser wavelength.
Diag Capable	Indicates if the transceiver is capable of doing diagnostics.
Connector Code	The vendor organizationally unique identifier field (OUI) contains the IEEE company identifier for the vendor.
Vendor OUI	The vendor-specific identifier field (OUI) contains the IEEE company identifier for the vendor.
Manufacture date	The manufacturing date of the hardware component in the mmddyyyy ASCII format.
Media	The media supported for the SFP.
Serial Number	The vendor serial number of the hardware component.
Part Number	The vendor part number contains ASCII characters, defining the vendor part number or product name.
Optical Compliance	Specifies the optical compliance code of the transceiver.
Link Length support	Specifies the link length support for the transceiver.
Transceiver Digital Diagnostic Monitoring (DDM)	Displays information for the transceiver Digital Diagnostic Monitoring (DDM), such as temperature and supply voltage.
Traffic Statistics	
Input/Output	When the collection of accounting and statistical data is enabled, then octet, packet, error, and utilization statistics are displayed.

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
Errors Input/Output	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol. For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.
Utilization Input/Output	The value computed as the average of the traffic observed over the configured interval, presented as a percentage of the maximum possible traffic.
Ethernet Statistics	See Table 76 .
Port Statistics	
Unicast Packets Input/Output	The number of packets, delivered by this sub-layer to a higher (sub-) layer, which were not addressed to a multicast or broadcast address at this sub-layer. The total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent.
Multicast Packets Input/Output	The number of packets, delivered by this sub-layer to a higher (sub-) layer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol, this includes both group and functional addresses. The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
Broadcast Packets Input/Output	The number of packets, delivered by this sub-layer to a higher (sub-) layer, which were addressed to a broadcast address at this sub-layer. The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
Discards Input/Output	The number of inbound packets chosen to be discarded to possibly free up buffer space.
Unknown Proto Discards Input/Output	For packet-oriented interfaces, the number of packets received through the interface which were discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing the number of transmission units received via the interface which were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter will always be 0. For ATM, this field displays cells discarded on an invalid vpi/vci. Unknown proto discards do not show up in the packet counts.
Ethernet-like Medium Statistics	See Table 77 .

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
Per Threshold MDA Discard Statistics	
Ingress Port Forwarding Engine Drop Reason Statistics	When any of the packet counters increments, it indicates that a packet was dropped in the ingress datapath or extracted to the control plane for further processing. The detailed per-reason drop statistics are available per-port for Ethernet ports (local on the 7450 ESS, 7750 SR, 7950 XRS, or Satellite client ports) and for PXC sub-ports. An aggregate forwarding engine drop counter (packet and byte) is available per-SAP.
IPv4 Header Error	The IPv4 packet header contains an error such as an IPv4 header checksum error, an invalid IP version number (not 4 or 6), or an incorrect Total Length field.
IPv4 Invalid Address	An error in the source (SA) or destination (DA) IPv4 address was detected. For example class D or class E IPv4 DAs, loopback SA (127.0.0.0/8), 0.0.0.0/8 DA, SA is a subnet broadcast or network address and cases where the IPv4 address is a multicast address but the Ethernet destination address is not RFC1112 compliant. RFC1112 checks are also carried out on IPIPE traffic.
IPv6 Header Error	The IPv6 packet header contains an error such as an incorrect Payload Length field or an IP version not equal to '6' when the Ethernet etype or PPP ID indicates it is IPv6.
IPv6 Invalid Address	An error in the source (SA) or destination (DA) IPv6 address was detected. For example, an unspecified IPv6 DA, an IPv6 multicast SA and cases where the IPv6 address is a multicast address but the Ethernet destination address is not RFC2464 compliant. RFC 2464 checks are also carried out on IPIPE traffic.
IP Route Blackholed	The destination IP address of the packet matches a black hole route.
ACL Filter Discards	The packet was dropped by a filter (ACL) on the line card (such as IP or MAC filter). Packets dropped by CPM filters or ESM antispoof filters do not increment this counter.
Unicast RPF Check Failed	The IP packet failed the unicast reverse path forwarding (uRPF) check.
BFD Spoof Check Failed	The received BFD packet either failed the TTL check (single hop BFD TTL should be 255) or failed the source IP address lookup of known sessions. This counter may increment when BFD sessions are first configured since one side may start sending before the other is completely programmed and ready.
Unicast MAC Destination Address Mismatch	The unicast destination MAC address is a null address or does not match any of the expected MAC addresses associated with the receiving interface. Packets with multicast or broadcast MAC addresses do not increment this counter.
Multicast MAC With Unicast Dest IP	The Ethernet destination (MAC) address is multicast but the IP address is unicast.
Unknown MAC Destination Address Discarded in VPLS	The MAC destination address lookup in the MAC FDB failed and the VPLS service is configured to discard packets with unknown destination MAC addresses.

Table 74 **Output Fields: show port <port-id> detail (Continued)**

Label	Description (Continued)
L2 Service MTU Exceeded	The length of the packet received on a SAP bound to a layer 2 service (such as VPLS or EPIPE) exceeded the configured MTU for the service.
Needs ICMP	The received packet requires the router to generate an ICMP message. For example, the IP packet TTL is expired or the destination host, network, or Enhanced Subscriber Management (ESM) subscriber is unreachable. Host unreachable can occur, among other reasons, if the destination address (e.g. 10.0.1.2) of a packet resolves to a loopback interface subnet (e.g. 10.0.1.1/24) but doesn't match the specific loopback address (e.g. 10.0.1.1). Destination unreachable can also occur, for example, in response to an IPv6 packet received by a router from a point-to-point link (a non-Ethernet link) destined to an address within a subnet assigned to that same link (other than one of the receiving router's own addresses) as described in RFC 4443.

Sample Output: show port <port id> optical detail

```

*A:ALU-1# show port 4/1/1 optical detail
=====
Optical Interface
=====
Transceiver Data

Transceiver Status : operational
Transceiver Type   : MSA-100GLH
Model Number      : 28-0089-XX
TX Laser Wavelength: 1558.172 nm
TX Laser Frequency : 192.400 THz
Laser Tunability   : fully-tunable
Present Channel    : 24
Configured Chann* : 24
50GHz Ch Min/Max  : 115/605
100GHz Ch Min/Max : 12/61
DAC Percent        : 50.00 %

RxDTV Adjust       : Enabled
Diag Capable       : yes
Number of Lanes    : 1
Connector Code     : LC
Vendor OUI         : 00:03:fa
Manufacture date   : 2012/07/16
Media              : Ethernet
Serial Number      : 122900645
Part Number        : AC100-201-00E
Optical Compliance : DWDM-TUN
Link Length support: 80km for SMF

=====
Transceiver Digital Diagnostic Monitoring (DDM)
=====
Value High Alarm  High Warn  Low Warn  Low Alarm
-----
Temperature (C)   +60.9    +80.0    +70.0    +0.0     -5.0
Supply Voltage (V) 12.07    13.00    12.60    11.40    11.00
=====

=====
Transceiver Lane Digital Diagnostic Monitoring (DDM)
=====
High Alarm  High Warn  Low Warn  Low Alarm
-----

```

Lane Temperature (C)	+75.0	+70.0	+20.0	+15.0
Lane Tx Bias Current (mA)	10.0	9.0	3.0	2.0
Lane Tx Output Power (dBm)	3.00	2.00	0.00	-1.00
Lane Rx Optical Pwr (avg dBm)	8.16	5.00	-20.00	-23.01

Lane ID	Temp(C)/Alm	Tx Bias(mA)/Alm	Tx Pwr(dBm)/Alm	Rx Pwr(dBm)/Alm
1	+48.4	5.1	0.99	-10.45

Coherent Optical Module

fg Tx Target Power: 1.00 dBm Present Rx Channel : 24
Cfg Rx LOS Thresh : -23.00 dBm Cfg Rx Channel : 24 (auto)

Disp Control Mode : automatic Sweep Start Disp : -25500 ps/nm
Cfg Dispersion : 0 ps/nm Sweep End Disp : 2000 ps/nm
CPR Window Size : 4 symbols

Cfg Alarms : modflt mod netrx nettx hosttx
Alarm Status :
Defect Points :

Rx Q Margin : 10.1 dB Chromatic Disp : 1 ps/nm
SNR X Polar : 19.7 dB Diff Group Delay : 0 ps
SNR Y Polar : 19.8 dB Pre-FEC BER : 0.000E+00

Module State : ready
Tx Turn-Up States : init laserTurnUp laserReadyOff laserReady
modulatorConverge outputPowerAdjust
Rx Turn-Up States : init laserReady waitForInput adcSignal opticalLock
demodLock

Coherent Optical Port Statistics (Elapsed Seconds: 80674)

Statistic	Current	Average	Minimum	Maximum
Rx BER	0.000E+00	2.323E-05	0.000E+00	4.646E-05
Rx SNR (dB)	19.6	10.0	0.0	20.1
Rx Chromatic Disp (ps/nm)	1	-18	-37	1
Rx Diff Group Delay (ps)	0	0	0	0
Rx Freq Offset (MHz)	38	-74	-347	200
Rx Q (dB)	16.6	8.3	0.0	16.6
Rx Power (dBm)	-10.44	-13.40	-99.00	-10.39
Tx Power (dBm)	0.98	-2.00	-99.00	1.01

Wavelength Tracker

Port Type : pluggable SFP VOA Present : yes
SFP VOA Type : fast
Serial Number : ALLU11--JS0100456

```

Part Number      : 21131722-0101

Power Control    : Enabled                WaveKey Status    : Disabled
Target Power     : -10.00 dBm             WaveKey 1           : 0
Measured Power   : -9.99 dBm             WaveKey 2           : 0

Cfg Alarms       : enc-fail enc-degr pwr-fail pwr-degr pwr-high pwr-low
                  : missing
Alarm Status      :

Maximum Power     : -2.60 dBm             Power Upper Margin : 7.39 dB
Minimum Power     : -22.00 dBm            Power Lower Margin : 12.01 dB
=====

```

Table 75 **Output Fields: show port <port-id> optical detail**

Label	Description
Transceiver Status	Status of the transceiver.
Transceiver Type	Type of the transceiver.
Model Number	The model number of the transceiver.
Transceiver Code	The code for the transmission media. <Q: is this field still available in any outputs?>
TX Laser Wavelength	Indicates the transceiver laser wavelength.
Present Channel	Specifies the present channel that the transceiver is on.
TX Laser Frequency	Indicates the transceiver laser frequency.
Configured Chann*	Specifies the channel that is configured for the transceiver.
Laser Tunability	Specifies the laser tune-ability of the transceiver.
50GHz Ch Min/Max	Specifies the 50 GHz channel minimum/maximum.
100GHz Ch Min/Max:	Specifies the 100 GHz channel minimum/maximum.
RxDTV Adjust	Specifies the Rx DTV adjust status.
DAC Percent	Specifies the DAC percentage of the transceiver.
Diag Capable	Indicates if the transceiver is capable of doing diagnostics.
Number of Lanes	Specifies the number of lanes of the transceiver.

Table 75 **Output Fields: show port <port-id> optical detail (Continued)**

Label	Description (Continued)
Connector Code	The vendor organizationally unique identifier field (OUI) contains the IEEE company identifier for the vendor.
Vendor OUI	Specifies the vendor OUI of the transceiver.
Manufacture date	The manufacturing date of the hardware component in the mmddyyyy ASCII format.
Media	The media supported for the SFP.
Serial Number	The vendor serial number of the hardware component.
Part Number	The vendor part number contains ASCII characters, defining the vendor part number or product name.
Optical Compliance	Specifies the optical compliance code of the transceiver.
Link Length support	Specifies the link length support for the transceiver.
Transceiver Digital Diagnostic Monitoring (DDM)	Displays information for the transceiver Digital Diagnostic Monitoring (DDM), such as temperature and supply voltage.
Transceiver Lane Digital Diagnostic Monitoring (DDM)	Displays information for the transceiver lane DDM, such as lane temperature and Tx bias current.
Coherent Optical Module	Displays information for the coherent optical module.
Coherent Optical Port Statistics	Displays statistics for the coherent optical port. For Coherent CFP, the Rx Power field displays per-channel power.
Wavelength Tracker	Displays wavelength tracker information.

Sample Output: show port <port-id> vport <vport-name>

```
*A:Bennet-Dut-A# show port 1/1/2 vport "vp1"
=====
Ethernet port 1/1/2 Access Egress vport
=====
VPort Name      : vp1
Description     : (Not Specified)
Sched Policy    : psp

Host-Matches
-----
Dest: dslam1
-----
```

```

=====
*A:Bennet-Dut-A#

*A:Bennet-Dut-A# show port 1/1/2 vport "vp1" associations
=====
Ethernet port 1/1/2 Access Egress vport
=====
-----
VPort "vp1"
-----
svc-id : 1
sap    : 1/1/2:1
subscr: s1
ip      : 10.1.1.2
mac     : 00:00:00:00:00:01  pppoe-sid: N/A
=====
*A:Bennet-Dut-A

*A:Bennet-Dut-A# show port 1/1/1 vport "abc" monitor-threshold
=====
Port 1/1/1 Vport "abc" Monitor Threshold Info
=====
Attribute                                Exceed Count Config Rate    Threshold Prcnt
-----
Agg-Eps                                0              212              32
Lvl-1                                  0             12323             89
Lvl-2                                  0             32132             32
Lvl-5                                  0              2323              4
Grp-01234567890123458746513513355656 0              2121              12
-----
Start Time      : 01/07/2015 16:53:16      End Time       : 01/07/2015 16:53:36
Total Samples :
=====
*A:Bennet-Dut-A

```

Note: If the Vport name is omitted, statistics for all Vports are displayed (bulk read). The statistics are displayed only for the levels, groups, and agg-eps for which the monitor-threshold is enabled. The output information filtering per level, group, or agg-eps is not embedded in the show commands natively. Instead, the output can be filtered with the match extensions for the show command. For example, show port 1/1/1 vport test monitor-threshold | match Lvl-1.

```

*A:sne# show port 1/1/4 vport statistics
=====
Port 1/1/4 Access Egress vport
=====
VPort Name      : vp1
Description     : (Not Specified)
Sched Policy    : portschedpoll
Rate Limit      : Max
Rate Modify     : disabled
Modify delta    : 0
Vport Queueing  Statistics

Last Cleared Time : N/A

```

```

                                Packets          Octets
Forwarded:                     0                0
Dropped :                      0                0
-----
Vport per Level Queueing Statistics
                                Packets          Octets
Level : 8
Forwarded:                     0                0
Dropped :                      0                0
Level : 7
Forwarded:                     0                0
Dropped :                      0                0
Level : 6
Forwarded:                     0                0
Dropped :                      0                0
Level : 5
Forwarded:                     0                0
Dropped :                      0                0
Level : 4
Forwarded:                     0                0
Dropped :                      0                0
Level : 3
Forwarded:                     0                0
Dropped :                      0                0
Level : 2
Forwarded:                     0                0
Dropped :                      0                0
Level : 1
Forwarded:                     0                0
Dropped :                      0                0

Host-Matches
-----
Dest: dslaml
-----
=====
*A:sne#

```

Sample Output: show port <port-id> detail (Excerpt Showing Ethernet Statistics)

```

A:ALA-48# show port 1/3/1 detail
=====
...
=====
Ethernet Statistics
=====
Broadcast Pckts :                42621  Drop Events      :                0
Multicast Pckts :                   0  CRC/Align Errors :                0
Undersize Pckts :                   0  Fragments        :                0
Oversize Pckts  :                   0  Jabbers          :                0
Collisions      :                   0

Octets          :                2727744
Packets         :                42621
Packets of 64 Octets :                42621
Packets of 65 to 127 Octets :                0
Packets of 128 to 255 Octets :                0
Packets of 256 to 511 Octets :                0

```

```

Packets of 512 to 1023 Octets :          0
Packets of 1024 to 1518 Octets :         0
Packets of 1519 or more Octets :         0
=====
Port Statistics
=====
                                     Input          Output
-----
Unicast Packets                      0              0
Multicast Packets                     0              0
Broadcast Packets                    42621             0
Discards                             0              0
Unknown Proto Discards                0
=====
...

```

Output Fields: show port <port-id> detail (Excerpt Showing Ethernet Statistics)

[Table 76](#) describes the output fields for the **show port detail** command (showing the Ethernet Statistics section).

Table 76 **Output Fields: show port <port-id> detail (Excerpt Showing Ethernet Statistics)**

Label	Description
Broadcast Pckts	The number of packets, delivered by this sub-layer to a higher (sub-) layer, which were addressed to a broadcast address at this sub-layer. The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
Drop Events	The total number of events in which packets were dropped by the probe due to lack of resources. Note that this number is not necessarily the number of packets dropped; it is just the number of times this condition has been detected.
Multicast Pckts	The number of packets, delivered by this sub-layer to a higher (sub-) layer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol, this includes both Group and Functional addresses. The total number of packets that higher-level protocols requested be transmitted, and which were addressed to a multicast address at this sub-layer, including those that were discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
CRC/Align Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
Undersize Pckts	The total number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.

Table 76 Output Fields: show port <port-id> detail (Excerpt Showing Ethernet Statistics)

Label	Description (Continued)
Fragments	The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
Oversize Pckts	The total number of packets received that were longer than can be accepted by the physical layer of that port (9900 octets excluding framing bits, but including FCS octets for GE ports) and were otherwise well formed.
Jabbers	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
Collisions	The best estimate of the total number of collisions on this Ethernet segment.
Octets	The total number of octets received.
Packets	The total number of packets received.

Sample Output: show port <port-id> detail (Excerpt Showing Ethernet-like Medium Statistics)

```
A:ALA-48# show port 1/3/1 detail
=====
...
=====
Ethernet-like Medium Statistics
=====
Alignment Errors :          0  Sngl Collisions :          0
FCS Errors       :          0  Mult Collisions :          0
SQE Test Errors  :          0  Late Collisions :          0
CSE              :          0  Excess Collisns :          0
Too long Frames  :          0  Int MAC Tx Errs :          0
Symbol Errors    :          0  Int MAC Rx Errs :          0
=====
A:ALA-48#
```

```
A:ALA-48# show port 1/2/1.sts192
=====
WAN Interface Sublayer Path Info
=====
Oper Status      : up
Signal Label     : 0x1a          Rx Signal Label      : 0x1a
Trace String     : Alcatel 7750 SR
Rx Trace Str(Hex) : 89 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Cfg Alarm        : pais plop prdi pplm prei puneq plcd
Alarm Status     :
=====
Port Statistics
```

```

=====
                                     Input
Output
-----
Unicast Packets                      367218143
5311
Multicast Packets                     0
0
Broadcast Packets                     0
0
Discards                             0
0
Unknown Proto Discards                0
=====
A:ALA-48#

A:ALA-48# show port 1/2/1.sts192 detail
=====
WAN Interface Sublayer Path Info
=====
Oper Status       : up
Signal Label      : 0x1a          Rx Signal Label      : 0x1a
Trace String      : Alcatel 7750 SR
Rx Trace Str(Hex) : 89 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Cfg Alarm         : pais plop prdi pplm prei puneq plcd
Alarm Status      :
=====
Sonet Path
=====
                                     Far End
-----
ES-P              1              0
SES-P             1              0
UAS-P             0              0
CV-P              10             0
=====
Transmit:
LOP-P             :              0
Fifo Error        :              96
Max Packet Error  :              0
Min Packet Error  :              0
LLP Packet Error  :              0
FIFO Underflow Error :          0

Receive:
LOP-P             :              0
AIS-P             :              1
RDI-P             :              0
PLM-P             :              0
LCD-P             :              0
Unequipped        :              0
Remote Error      :              6
Parity Error      :             10
Fifo Error        :              0
Max Pkt Error     :              0
Min Pkt Error     :              0

```

```

FCS Error          :          0
Packet Abort Error :          0
Addr Ctrl Invalid  :          0
=====
Port Statistics
=====
Input
Output
-----
Unicast Packets          369758853
5312
Multicast Packets          0
0
Broadcast Packets          0
0
Discards          0
0
Unknown Proto Discards          0
=====
A:ALA-48#

```

Output Fields: show port <port-id> detail (Excerpt Showing Ethernet-like Medium Statistics Output)

[Table 77](#) describes the output fields for the **show port detail** command (showing the Ethernet-like Medium Statistics section).

Table 77 Output Fields: show port detail (Excerpt Showing Ethernet-like Medium Statistics)

Label	Description
Alignment Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets.
Sngl Collisions	The number of frames that are involved in a single collision, and are subsequently transmitted successfully.
FCS Errors	The number of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check.
Mult Collisions	The number of frames that are involved in more than one collision and are subsequently transmitted successfully.
SQE Errors	The number of times that the SQE TEST ERROR is received on a particular interface.
Late Collisions	The number of times that a collision is detected on a particular interface later than one slotTime into the transmission of a packet.
CSE	The number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface.

Table 77 **Output Fields: show port detail (Excerpt Showing Ethernet-like Medium Statistics)**

Label	Description (Continued)
Excess Collisns	The number of frames for which transmission on a particular interface fails due to excessive collisions.
Too long Frames	The number of frames received on a particular interface that exceed the maximum permitted frame size.
Int MAC Tx Errs	The number of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error.
Symbol Errors	For an interface operating at 100 Mb/s, the number of times there was an invalid data symbol when a valid carrier was present.
Int MAC Rx Errs	The number of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error.
In Pause Frames	The number of In Pause frames.
Out Pause Frames	The number of Out Pause frames.

Sample Output: show port <port-id> (Showing Channelized Ports)

```

A:ALA-7# show port 7/1/1.ds0grp-1.1
=====
TDM DS0 Chan Group
=====
Description      : DS3
Interface        : 7/1/1.ds0grp-1.1
TimeSlots       : 1
Speed           : 64
Admin Status    : up
Last State Change : 2007/04/11 01:14:37
CRC             : 16
Oper status     : down
Chan-Grp IfIndex : 656441433

Configured mode  : access
Admin MTU       : 1522
Physical Link    : No
Encap Type      : bcp-null
Oper MTU        : 1522
=====
Port Statistics
=====
                        Input      Output
-----
Unicast Packets      0          0
Multicast Packets    0          0
Broadcast Packets    0          0
Discards             0          0
Unknown Proto Discards 0          0
=====
A:ALA-7#

A:ALA-7# show port 7/1/1.ds0grp-1.1 detail

```

```

=====
TDM DS0 Chan Group
=====
Description      : DS3
Interface        : 7/1/1.ds0grp-1.1
TimeSlots       : 1
Speed           : 64                      CRC                : 16
Admin Status    : up                      Oper status        : down
Last State Change : 04/11/2007 01:14:37   Chan-Grp IfIndex    : 656441433

Configured mode  : access                  Encap Type          : bcp-null
Admin MTU        : 1522                    Oper MTU            : 1522
Physical Link    : No
=====
Port Statistics
=====
                                     Input                Output
-----
Unicast Packets          0                        0
Multicast Packets        0                        0
Broadcast Packets        0                        0
Discards                 0                        0
Unknown Proto Discards   0
=====
A:ALA-7#

ALA-12# show port 7/1/1.1.1
=====
TDM DS0 Chan Group
=====
Description      : DS0GRP
Interface        : 3/1/1.1.1
TimeSlots       : 1
Speed           : 64                      CRC                : 16
Admin Status    : up                      Oper status        : down
Last State Change : 04/11/2007 06:54:28   Chan-Grp IfIndex    : 589332542

Configured mode  : access                  Encap Type          : bcp-null
Admin MTU        : 1518                    Oper MTU            : 1518
Physical Link    : No                      Bundle Number       : none
Idle Cycle Flags : flags                   Load-balance-algo   : default
=====
Traffic Statistics
=====
                                     Input                Output
-----
Octets                 0                        0
Packets                0                        0
Errors                 0                        0
=====
Port Statistics
=====
                                     Input                Output
-----
Packets                0                        0
Discards               0                        0
Unknown Proto Discards 0
=====

```

ALA-12#

A:ALA-48# show port 3/1/3.e3

=====

TDM Interface

```

=====
Description      : E3
Interface        : 3/1/3.e3
Type             : e3
Admin Status     : up
Physical Link    : No
Last State Change : 04/11/2007 06:54:28
Framing          : g751
Oper status      : down
Clock Source     : loop-timed
Port IfIndex     : 589398019

Configured mode  : access
Admin MTU        : 1518
CRC              : 16
Idle Cycle Flags : flags
FEAC Loop Respond : Disabled
BERT Pattern     : none
Err Insertion Rate : 0
BERT Status      : idle
Encap Type       : bcp-null
Oper MTU         : 1518
Channelized      : none
Loopback         : none
In FEAC Loop     : No
BERT Duration    : N/A
Last BERT Synched : 0 Seconds
Last BERT Errors  : 0
Last BERT Total Bits : 0

Cfg Alarm        : ais los
Alarm Status     :
Subrate Mode     : none

```

MDL Transmit : none

Local MDL Information

```

-----
EIC      :      LIC      :
FIC      :      Unit     :
PFI      :
Idle Signal Port :
Test Signal Gen  :

```

Far End MDL Information

```

-----
EIC      :      LIC      :
FIC      :      Unit     :
PFI      :
Idle Signal Port :
Test Signal Gen  :

```

=====

Traffic Statistics

```

=====
                                Input      Output
-----
Octets                        0            0
Packets                       0            0
Errors                        0            0

```

=====

Port Statistics

```

=====
                                Input      Output
-----
Packets                       0            0
Discards                      0            0

```

```
Unknown Proto Discards          0
=====
A:ALA-48#
```

Output Fields: show port <port-id> (Showing Channelized Ports)

Table 78 describes the output fields for the **show port** command for a channelized port.

Table 78 Output Fields: show port <port-id> (Showing Channelized Ports)

Label	Description
Description	A text description of the port.
Interface	The port ID.
TimeSlots	Specifies the DS0 timeslot used in the T1/E1 channel-group.
Speed	Indicates the speed of the DS0 channels used in the associated channel-group.
CRC	Indicates the precision of the cyclic redundancy check. 16 — A 16-bit CRC calculation. 32 — A 32-bit CRC calculation. 32-bit CRC increases the error detection ability, but it also adds some performance overhead.
Admin Status	Up — The port is administratively up. Down — The administratively down.
Oper Status	Up — The port is operationally up. Down — The port is operationally down.
Last State Change	Displays the last time the operational status of the port changed state.
Chan-Grp IfIndex	Displays the channel group's interface index number which reflects its initialization sequence.
Configured Mode	network — The port is configured for transport network use. access — The port is configured for service access. Channelized ports are always access ports.
Encap Type	The type of encapsulation protocol configured on this port's network interface. bcp-null — Indicates that BCP is used as the NCP control protocol. dot1q — Indicates that ingress frames carry 802.1Q tags where each tag signifies a different service.
Oper MTU	The negotiated size of the largest packet which can be sent on the channel, specified in octets. For channels that are used to transmit network datagrams, this is the size of the largest network datagram that can be sent.
Physical Link	Indicates whether or not the port has a physical link.

Table 78 **Output Fields: show port <port-id> (Showing Channelized Ports) (Continued)**

Label	Description (Continued)
Unicast Packets Input/Output	The number of packets, delivered by this sub-layer to a higher (sub-) layer, which were not addressed to a multicast or broadcast address at this sub-layer.

Sample Output: show port <port-id> associations

```

A:ALA-1# show port 1/1/6 associations
=====
Interface Table
=====
Router/ServiceId          Name                      Encap Val
-----
Router: Base              if1000                   1000
Router: Base              if2000                   2000
-----
Interfaces
=====
A;ALA-1#

```

Output Fields: show port <port-id> associations

[Table 79](#) describes the output fields for the **show port <port-id> associations** command.

Table 79 **Output Fields: show port <port-id> associations**

Label	Description
Svc ID	The service identifier.
Name	The name of the IP interface.
Encap Value	The dot1q or qinq encapsulation value on the port for this IP interface.

Sample Output: show port <port-id> frame-relay

```

A:ALA-49>config>port# show port 8/1/2 frame-relay
=====
Frame Relay Info for 8/1/2
=====
Mode                : dte          LMI Type           : itu
FR Interface Status  : fault
N391 DTE             : 6           N392 DCE           : 3
N392 DTE             : 3           N393 DCE           : 4
N393 DTE             : 4           T392 DCE           : 15
T391 DTE             : 10
Tx Status Enquiry    : 0           Rx Status Enquiry   : 0
Rx Status Messages   : 0           Tx Status Messages  : 0
Status Message Timeouts : 0       Status Enquiry Timeouts : 0
Discarded Messages   : 0           Inv. RxSeqNum Messages : 0

```



```
=====
A:ALA-49>config>port#
```

Output Fields: show port <port-id> frame-relay

Table 80 describes the output fields for the **show port <port-id> frame-relay** command.

Table 80 Output Fields: show port <port-id> frame-relay

Label	Description
Mode	Displays the mode of the interface. It can be set as Data terminal equipment (dte) or Data circuit-terminating equipment (DCE).
LMI Type	Displays the LMI type.
FR Interface Status	Displays the status of the Frame Relay interface as determined by the performance of the DLCMI. If no DLCMI is running, the Frame Relay interface will stay in the running state indefinitely.

Sample Output: show port <port-id> otu detail

```
A:ALA-49>config>port# show port 3/2/1 otu detail
=====
OTU Interface
=====
OTU Status          : Enabled          FEC Mode           : enhanced
                   :                   Data Rate          : 11.049 Gb/s
Cfg Alarms          : loc los lof lom otu-ber-sf otu-bdi fec-sf
Alarm Status        :
SF/SD Method        : FEC              SF Threshold       : 1E-5
                   :                   SD Threshold       : 1E-7

SM-TTI Tx (auto)    : ALA-49:3/2/1/C17
SM-TTI Rx           : (Not Specified)
=====
OTU Statistics
=====
Statistics                      Count
-----
FEC Corrected 0s                0
FEC Corrected 1s                0
FEC Unrectable Sub-rows         0
FEC SES                         0
SM BIP8                         0
SM BEI                         0
SM SES                         0
PM BIP8                         0
PM BEI                         0
PM SES                         0
=====
```

Output Fields: show port <port-id> otu detail

[Table 81](#) describes the output fields for the **show port <port-id> otu detail** command.

Table 81 Output Fields: show port <port-id> otu detail

Label	Description
OTU Status	Status of the OTU (Optical Transport Unit): enabled or disabled. When OTU is enabled, an additional layer of framing encapsulates an MDA's natively programmed mode of operation, 10-Gigabit Ethernet LAN or WAN, adding SONET-Like Framing with FEC (Forward Error Correction). When OTU is disabled, the MDA operates in a 10-Gigabit Ethernet LAN or WAN as per Ethernet provisioning.
FEC Mode	Type of FEC (Forward Error Correction) in effect: g709, enhanced or disabled. When g709 is selected, the standard FEC method is used. When enhanced is selected, a proprietary FEC algorithm is used that extends optical reach in long haul applications. When disabled the bytes that are reserved for FEC in the OTU frame are transmitted as zeros and the FEC decoder is bypassed, but OTU framing is still in effect.
Data Rate	This indicates the data rate at which the port is operating. When OTU is encapsulating 10-Gigabit Ethernet WAN, the data rate is 10.709 Gb/s, the G.709 standard OTU2 data rate. When OTU is encapsulating 10-Gigabit Ethernet LAN, the data rate is either 11.049 Gb/s or 11.096 Gb/s, depending on the otu2-lan-data-rate configuration parameter of the port's OTU parameters. These data rates (11.049 Gb/s and 11.096 Gb/s) are considered OTU2e data rates that are non-standard or over-clocked with respect to G.709, but have become widely used in optical networking to transport un-altered 10-Gigabit Ethernet LAN payloads.
Cfg Alarms and Alarm Status	This indicates the alarms that shall be reported when raised or cleared. Alarms that are not in this list will not be reported when they are raised or cleared but will appear in the Alarm Status.
SF/SD Method	This indicates the selected method for declaring the SF (Signal Fail) or SD (Signal Degrade) alarm. When BIP8 is selected, the error rate of SM-BIP8 errors in the OTU frames declares SF or SD (This is very similar to SONET SF/SD which uses a rate of B2 errors). When FEC is selected, the rate of corrected bits declares SF or SD. This effectively indicates that the link would be degraded (SD) or failed (SF) if FEC was disabled and gives the user an early warning that the link is degrading or is about to fail.
SF Threshold	This is the configured error rate threshold at which the SF (Signal Fail) alarm will be raised.
SD Threshold	This is the configured error rate threshold at which the SD (Signal Degrade) alarm will be raised.
SM-TTI Tx (<mode>)	This is the configured SM-TTI (Section Monitor Trail Trace Identifier) to be transmitted by this port in the OTU overhead bytes. The modes are auto, string, or bytes. In the auto and string modes, a printable character string will be displayed. In bytes mode, up to 64 hex bytes will be displayed.

Table 81 **Output Fields: show port <port-id> otu detail (Continued)**

Label	Description (Continued)
SM-TTI Rx	This is the SM-TTI (Section Monitor Trail Trace Identifier received by this port. When the received TTI is a printable string of characters, it will be displayed as a text string. When the received TTI contains one or more non-printable characters, it will be displayed at a sequence of 64 hex bytes. When the received TTI is all zeros, the string "Not Specified" will be displayed.
FEC Corrected 0s	Displays the number of bits that were received as 0s but corrected to 1s.
FEC Corrected 1s	Number of bits that were received as 1s but corrected to 0s.
FEC Uncorrectable Sub-Rows	The number of sub-rows that were not corrected because too many errors were detected.
FEC SES	The number of severely errored seconds where the number of uncorrectable sub-rows was greater than 15% of the maximum.
SM BIP8	The number of detected BIP-8 errors in the section monitor overhead.
SM BEI	The number of backward error indications received from the far end in the section monitor overhead.
SM SES	Section monitor severely errored seconds where the number of SM-BIP8 was greater than 15% of the maximum.
PM BIP8	The number of detected BIP-8 errors in the path monitor overhead.
PM BEI	The number of backward error indications received from the far end in the path monitor overhead.
PM SES	Section monitor severely errored seconds where the number of PM-BIP8 was greater than 15% of the maximum.

Sample Output: show port <port-id> ppp

```
A:SR-007# show port 1/1/1.1.1.1.1 ppp
=====
PPP Protocols for 1/1/1.1.1.1
=====
Protocol   State           Last Change           Restart Count   Last Cleared
-----
lcp        opened          03/28/2007 13:06:28           7       03/28/2007 12:12:11
ipcp       initial         03/28/2007 11:39:45           0       03/28/2007 12:12:11
mplscp     initial         03/28/2007 11:39:45           0       03/28/2007 12:12:11
bcp        initial         03/28/2007 11:39:45           0       03/28/2007 12:12:11
osicp      opened          03/28/2007 13:06:28          12       03/28/2007 12:12:11
ipv6cp     opened          03/28/2007 13:06:28           7       03/28/2007 12:12:11
=====
PPP Statistics
=====
Local Mac address   : 00:03:fa:0e:76:e2   Remote Mac address : 00:00:00:00:00:00
Local Magic Number  : 0x7e9a9       Remote Magic Number: 0x7e18b
```

```

Local IPv4 address : 0.0.0.0           Remote IPv4 address: 0.0.0.0
Local IPv6 address : FE80::203:FAFF:FE81:5888
Remote IPv6 address: FE80::203:FAFF:FE1A:7EE2
Line Monitor Method: keepalive
Keepalive statistics
Request interval   : 10           Threshold exceeded : 0
Drop Count        : 3           In packets          : 332
Time to link drop : 00h00m30s    Out packets         : 332
Last cleared time  : 03/28/2007 12:12:11
=====
A:SR-007#

```

```

A:SR-007# show port 1/1/3 ppp detail
=====
PPP Protocols for 1/1/3.sts12
=====
Protocol  State          Last Change          Restart Count    Last Cleared
-----
lcp       initial          04/11/2007 10:56:11      0      04/11/2007 10:56:11
ipcp      initial          04/11/2007 10:56:11      0      04/11/2007 10:56:11
mplscp    initial          04/11/2007 10:56:11      0      04/11/2007 10:56:11
bcp       initial          04/11/2007 10:56:11      0      04/11/2007 10:56:11
osicp     initial          04/11/2007 10:56:11      0      04/11/2007 10:56:11
=====
PPP Statistics
=====
Local IP address   : 0.0.0.0           Remote IP address   : 0.0.0.0
Local Mac address  : 00:00:00:00:00:00    Remote Mac address  : 00:00:00:00:00:00
Local Magic Number : 0x0               Remote Magic Number : 0x0
Line Monitor Method: keepalive
Keepalive statistics
Request interval   : 10           Threshold exceeded : 0
Drop Count        : 3           In packets          : 0
Time to link drop : 00h00m30s    Out packets         : 0
Last cleared time  : 04/11/2007 10:56:11
=====
A:SR-007#

```

The following output displays an example of a PPP link inside a multilink-bundle group:

```

*A:top_SR7# show port 3/1/4.1.2.3.1 ppp
=====
LCP Protocol for 3/1/4.1.2.3.1
=====
Protocol  State          Last Change          Restart Count    Last Cleared
-----
lcp       opened          05/22/2008 07:46:18      1      05/22/2008 06:28:48
=====
Keepalive statistics
Request interval   : 4           Threshold exceeded : 0
Drop Count        : 3           In packets          : 46404
Time to link drop : 00h00m12s    Out packets         : 46404
Last cleared time  : 05/22/2008 06:28:48

```

```

PPP Header Compression
  ACFC          : Enabled      PFC          : Enabled
=====
*A:top_SR7#

The following output displays an example of a standalone PPP link:

*A:top_SR7# show port 3/1/
4.1.2.3.1 ppp
=====
PPP Protocols for 3/1/4.1.2.3.1
=====
Protocol  State          Last Change          Restart Count    Last Cleared
-----
lcp       initial          05/24/2008 11:25:23      1      05/22/2008 06:28:48
ipcp      initial          05/22/2008 06:28:48      0      05/22/2008 06:28:48
mplscp    initial          05/22/2008 06:28:48      0      05/22/2008 06:28:48
bcp       initial          05/22/2008 06:28:48      0      05/22/2008 06:28:48
osicp     initial          05/22/2008 06:28:48      0      05/22/2008 06:28:48
ipv6cp    initial          05/22/2008 06:28:48      0      05/22/2008 06:28:48
=====
PPP Statistics
=====
Local Mac address   : 00:16:4d:8f:d3:57  Remote Mac address :
Local Magic Number  : 0x0          Remote Magic Number: 0x0
Local IPv4 address  : 0.0.0.0       Remote IPv4 address: 0.0.0.0
Local IPv6 address  : ::
Remote IPv6 address: ::

Line Monitor Method: keepalive

Keepalive statistics

Request interval   : 4          Threshold exceeded : 0
Drop Count        : 3          In packets         : 46418
Time to link drop  : 00h00m12s  Out packets        : 46418
Last cleared time  : 05/22/2008 06:28:48

PPP Header Compression
  ACFC          : Enabled      PFC          : Enabled
=====
*A:top_SR7#

```

Output Fields: show port <port-id> ppp

[Table 82](#) describes the output fields for the **show port <port-id> ppp** command.

Table 82 Output Fields: show port <port-id> ppp

Label	Description
Protocol	Displays the applicable protocols for the specified port.

Table 82 **Output Fields: show port <port-id> ppp (Continued)**

Label	Description (Continued)
State	Displays the current status of a PPP link. Values include initial, starting, closed, stopped, closing, stopping, requestSent, ackReceived, ackSent, opened.
Last Change	Displays the last time the PPP link state changed.
Restart Count	Displays the number of times that this Control Protocol has reached the 'open' state.
Last Cleared	Displays the date and time the restart count was set to zero.
Local IP address	Displays the IP address assigned the local end of the PPP link.
Remote IP address	Displays the IP address of the remote end of the PPP link.
Local Mac address	Displays the MAC address assigned the local end of the PPP link.
Remote Mac address	The Ethernet MAC address sent by the remote end of the PPP link.
Local Magic Number	Displays the local magic number to be sent to the peer. The magic number provides a method to detect loopbacks. If the value of the local magic number is the same as the value of remote magic number, then it is possible that the link might be looped back. If the two magic numbers do not match, then the link is not looped back.
Remote Magic Number	Displays the magic number sent by the peer. If the value of remote magic number is the same as the value of the local magic number, then it is possible that the link might be looped back. If the two magic numbers do not match, then the link is not looped back.
Line Monitor Method	The type of line monitoring packets being sent and received on this PPP link.
Request Interval	The time interval in seconds at which keepalive requests are issued.
Threshold exceeded	Displays the number of times that the drop count was reached.
Drop Count	Displays the number of keepalive or LQR messages that were missed before the line was brought down.
In Packets	Displays the number of echo-reply packets received.
Time to link drop	Displays the time remaining before the link will be declared dropped if a keepalive echo reply packet is not received.
Out packets	Displays the number of echo-request packets sent.
Last cleared time	Displays the time since the last clear.
ACFC	Indicates whether Address and Control Field PPP Header Compression is enabled.
PFC	Indicates whether Protocol Field PPP Header Compression is enabled.

Sample Output: show port <port-id> atm

```

ALA-1# show port 9/1/2 atm
=====
ATM Info for 9/1/2
=====
Cell Mode           : UNI
Configured VCs      : 1                      Max Supported VCs : 2000
Interface Oper Status : lower layer down   Number OCD Events : 0
TC Alarm State      : LCD Failure
Last Unknown VPI/VCI : none
=====
ATM Bandwidth Info
=====

```

	kbps	%		kbps	%
Ingress CBR	: 0	0%	Egress CBR	: 0	0%
Ingress RT-VBR	: 0	0%	Egress RT-VBR	: 0	0%
Ingress NRT-VBR	: 4000	1%	Egress NRT-VBR	: 0	0%
Ingress UBR	: 0	0%	Egress UBR	: 0	0%
Ingress Total	: 4000	1%	Egress Total	: 0	0%
ATM Link Bandwidth : 599041 kbps					
Shaped Bandwidth : 0 kbps					

```

=====
ALA-1#

*A:ALA-48# show port 3/2/1 atm detail
=====
ATM Info for 3/2/1
=====
Cell Mode           : UNI                      Mapping           : Direct
Configured VCs      : 0                      Configured VPs     : 0
Configured VTs      : 0                      Configured IFCs    : 0
Configured minimum VPI: 0
Last Unknown VPI/VCI : none
=====
TC Sublayer Information
=====
TC Alarm State      : LCD Failure           Number OCD Events : 0
HEC Errors (Dropped) : 0                  HEC Errors (Fixed) : 0
=====
ATM Bandwidth Info
=====

```

	kbps	%		kbps	%
Ingress CBR	: 0	0%	Egress CBR	: 0	0%
Ingress RT-VBR	: 0	0%	Egress RT-VBR	: 0	0%
Ingress NRT-VBR	: 0	0%	Egress NRT-VBR	: 0	0%
Ingress UBR	: 0	0%	Egress UBR	: 0	0%
Ingress Total	: 0	0%	Egress Total	: 0	0%
ATM Link Bandwidth : 599041 kbps					
Shaped Bandwidth : 0 kbps					

```

=====
ATM Statistics
=====

```

	Input	Output
Octets	0	0

```

Cells                                0                                0
Unknown VPI/VCI Cells                0
=====
AAL-5 Packet Statistics
=====
                                Input                                Output
-----
Packets                             0                                0
Dropped Packets                     0                                0
CRC-32 Errors                       0
=====
*A:ALA-48#

```

B:Dut-D# show port 2/2/1.1.1.1 atm

=====

ATM Info for 2/2/1.1.1.1

```

=====
Cell Mode          : UNI          Mapping          : Direct
Configured VCs     : 16           Configured VPs  : 0
Configured VTs     : 0           Configured IFCs : 0
Configured minimum VPI: 0
Last Unknown VPI/VCI : none
=====

```

TC Sublayer Information

```

=====
TC Alarm State      : No Alarm      Number OCD Events : 0
HEC Errors (Dropped) : 0          HEC Errors (Fixed) : 0
=====

```

ATM Bandwidth Info

```

=====
                                kbps      %                                kbps      %
-----
Ingress CBR        : 0          0%      Egress CBR        : 0          0%
Ingress RT-VBR     : 0          0%      Egress RT-VBR     : 0          0%
Ingress NRT-VBR    : 0          0%      Egress NRT-VBR    : 0          0%
Ingress UBR        : 0          0%      Egress UBR        : 0          0%
-----
Ingress Total      : 0          0%      Egress Total      : 0          0%
ATM Link Bandwidth : 1920 kbps
Shaped Bandwidth   : 0 kbps
=====

```

B:Dut-D#

B:Dut-D# show port 2/2/1.1.1.1 atm detail

=====

ATM Info for 2/2/1.1.1.1

```

=====
Cell Mode          : UNI          Mapping          : Direct
Configured VCs     : 16           Configured VPs  : 0
Configured VTs     : 0           Configured IFCs : 0
Configured minimum VPI: 0
Last Unknown VPI/VCI : none
=====

```

TC Sublayer Information

```

=====
TC Alarm State      : No Alarm      Number OCD Events : 0
HEC Errors (Dropped) : 0          HEC Errors (Fixed) : 0
=====

```



```

=====
ATM Bandwidth Info
=====
              kbps      %              kbps      %
-----
Ingress CBR      : 0      0%      Egress CBR      : 0      0%
Ingress RT-VBR   : 0      0%      Egress RT-VBR   : 0      0%
Ingress NRT-VBR  : 0      0%      Egress NRT-VBR  : 0      0%
Ingress UBR      : 0      0%      Egress UBR      : 0      0%
-----
Ingress Total    : 0      0%      Egress Total    : 0      0%
ATM Link Bandwidth : 1920 kbps
Shaped Bandwidth  : 0 kbps
=====
ATM Statistics
=====
                          Input                          Output
-----
Octets                        228425945553                228453511542
Cells                        4309923501                  4310443614
Unknown VPI/VCI Cells      4294967295
=====
AAL-5 Packet Statistics
=====
                          Input                          Output
-----
Packets                      4302445396                4302705455
Dropped Packets                      0                      0
CRC-32 Errors                      0
=====
A:timetra-sim110#

```

B:Dut-D# show port 2/1/1 atm connections

ATM Connections, Port 2/1/1

```

=====
              Owner  Type    Ing.TD  Egr.TD  Adm  OAM      Opr
-----
cp.1      SAP      CP       1       1     up    up       up
10/10     SAP      PVC      1       1     up    up       up
10/11     SAP      PVC      1       1     up    up       up
20/20     SAP      PVC      1       1     up    up       up
=====

```

B:Dut-D#

Sample Output: show port <port-id> atm cp

A:SR12# show port 3/1/1 atm cp

ATM CPs, Port 3/1/1

```

=====
CP      Owner  Type    Ing.TD  Egr.TD  Adm  OAM      Opr
-----
5       SAP      CP       1       1
9       SAP      CP       1       1
=====

```

A:SR12#

```

A:SR12# show port 3/1/1 atm cp detail
=====
ATM CP, Port 3/1/1
=====
CP          Owner  Type    Ing.TD  Egr.TD  Adm  OAM      Opr
-----
5           SAP    CP      1       1
=====
ATM CP Statistics
=====
Input              Output
-----
Octets                                284958830153      15950085
Cells                                5376581701       300945
Dropped CLP=0 Cells                                90           200
Dropped Cells (CLP=0+1)                        290
Tagged Cells                                780
=====
ATM CP, Port 3/1/1
=====
CP          Owner  Type    Ing.TD  Egr.TD  Adm  OAM      Opr
-----
9           SAP    CP      1       1
=====
ATM CP Statistics
=====
Input              Output
-----
Octets                                284958830153      15950085
Cells                                5376581701       300945
Dropped CLP=0 Cells                                90           200
Dropped Cells (CLP=0+1)                        290
Tagged Cells                                780
=====
A:SR12#

A:SR12# show port 3/1/1 atm cp 9
=====
ATM CP
=====
Port Id      : 3/1/1          CP      : 9
Owner        : SAP          Endpoint Type : CP
Ing. Td Idx  : 1           Egr. Td Idx  : 1
=====
A:SR12#

A:SR12# show port 3/1/1 atm cp 9 detail
=====
ATM CP
=====
Port Id      : 3/1/1          CP      : 9
Owner        : SAP          Endpoint Type : CP
Ing. Td Idx  : 1           Egr. Td Idx  : 1
=====

```

```

=====
ATM CP Statistics==
=====
Input                Output
-----
Octets                284958830153          15950085
Cells                5376581701           300945
Dropped CLP=0 Cells              90           200
Dropped Cells (CLP=0+1)         290
Tagged Cells                780
=====
A:SR12#

B:Dut-D# show connection-profile
=====
Connection Profile Summary Information
=====
CP Index  Number of
          Members
-----
1          3
=====
B:Dut-D#

B:Dut-D# show connection-profile 1
=====
Connection Profile 1 Information
=====
Description : My Connection Profile
Last Change : 09/11/2010 13:37:32

-----
VPI/VCI
-----
10/10
10/11
20/20
=====
B:Dut-D#

```

Output Fields: show port <port-id> atm

[Table 83](#) describes the output fields for the **show port <port-id> atm** command.

Table 83 **Output Fields: show port <port-id> atm**

Label	Description
Cell Mode	Displays the cell format (UNI or NNI) that is used on the ATM interface.
Configured VCs	Displays the number of configured VCs.
Max Supported VCs	Indicates the maximum number of ATM VPCs that can be configured on this MDA.

Table 83 **Output Fields: show port <port-id> atm (Continued)**

Label	Description (Continued)
Interface Oper Status	Indicates the status of the ATM interface. If the SONET-PATH layer and TC sublayer are operationally up, the ATM Interface is considered up. If the SONET-PATH layer and/or TC SubLayer is down, the ATM Interface is set to lowerLayerDown.
Number OCD Events	Displays the number of times the Out of Cell Delineation (OCD) events occurred.
TC Alarm State	Displays notifications that are generated when the ATM interface indicates that the TC sublayer is currently in the Loss of Cell Delineation (LCD) defect maintenance state or when the TC sublayer is currently not in the Loss of Cell Delineation (LCD) defect maintenance state.
Last Unknown VPI/VCI	Indicates the last unknown VPI/VCI that was received on this interface.
Ingress CBR	Indicates the total CBR bandwidth consumed on this interface in the ingress direction.
Egress CBR	Indicates the total CBR bandwidth consumed on this interface in the egress direction.
Ingress RT-VBR	Indicates the total real-time variable bit rate (rt-VBR) bandwidth consumed on this interface in the ingress direction.
Egress RT-VBR	Indicates the total real-time variable bit rate (rt-VBR) bandwidth consumed on this interface in the egress direction.
Ingress NRT-VBR	Indicates the total non-real-time variable bit rate (nrt-VBR) bandwidth consumed on this interface in the ingress direction.
Egress NRT-VBR	Indicates the total non-real-time variable bit rate (nrt-VBR) bandwidth consumed on this interface in the egress direction.
Ingress UBR	Indicates the total unspecified bit rate (UBR) bandwidth consumed on this interface in the ingress direction.
Egress UBR	Indicates the total unspecified bit rate (UBR) bandwidth consumed on this interface in the egress direction.
Ingress Total	Indicates the number of valid ATM cells received by the ATM interface including both CLP=0 and CLP=1 cells. If traffic policing is implemented, then cells are counted prior to the application of traffic policing.
ATM Link Bandwidth	Indicates the total ATM link bandwidth accepted on this interface.
Shaped Bandwidth	Indicates the total shaped bandwidth consumed on this interface in the egress direction.
HEC Errors (Dropped)	Indicates the number of cells with uncorrectable HEC errors on this interface.
HEC Errors (Fixed)	Indicates the number of fixed HEC errors on this interface.

Sample Output: show port <port-id> atm pvc detail

```
ALA-1# show port 9/1/2 atm pvc 0/500 detail
=====
ATM Endpoint
=====
Port Id       : 9/1/2          VPI/VCI       : 0/500
Admin State   : up            Oper state     : down
OAM State     : ETE-AIS       Encap Type     : llc
Owner         : SAP           AAL Type      : AAL-5
Endpoint Type : PVC           Cast Type      : P2P
Ing. Td Idx   : 5             Egr. Td Idx   : 3
Last Changed  : 02/14/2007 14:15:12
=====
ATM Statistics
=====
                                     Input          Output
-----
Octets                0                  0
Cells                 0                  0
=====
AAL-5 Packet Statistics
=====
                                     Input          Output
-----
Packets                0                  0
Dropped Packets        0                  0
CRC-32 Errors          0
Reassembly Timeouts    0
Over Sized SDUs        0
=====
ATM OAM Statistics
=====
                                     Input          Output
-----
AIS                    0                  0
RDI                    0                  0
Loopback               0                  0
CRC-10 Errors          0
Other                  0
=====
A:ALA-1#

A:ALA-228# show port 2/2/1.1.1.1 atm pvc 0/100 detail
=====
ATM PVC
=====
Port Id       : 2/2/1.1.1.1    VPI/VCI       : 0/100
Admin State   : up            Oper state     : up
OAM State     : up            Encap Type     : llc
Owner         : SAP           AAL Type      : AAL-5
Endpoint Type : PVC           Cast Type      : P2P
Ing. Td Idx   : 101           Egr. Td Idx   : 201
Last Changed  : 06/15/2007 22:09:11 ILMI Vpi/Vci Range : n/a
=====
ATM Statistics
=====
                                     Input          Output
-----
```

```

-----
Octets                    57173273                58892699
Cells                    1078741                 1111183
=====
AAL-5 Packet Statistics
=====
                                Input                Output
-----
Packets                    539382                555603
Dropped Packets              0                      0
CRC-32 Errors                0                      0
Reassembly Timeouts          0                      0
Over Sized SDUs              0                      0
=====
ATM OAM Statistics
=====
                                Input                Output
-----
AIS                          0                      0
RDI                          0                      1
Loopback                     0                      0
CRC-10 Errors                0                      0
Other                        0                      0
=====
A:ALA-228#

```

Output Fields: show port <port-id> atm pvc detail

[Table 84](#) describes the output fields for the **show port <port-id> atm pvc detail** command.

Table 84 **Output Fields: show port <port-id> atm pvc detail**

Label	Description
Port Id	Configures or displays the port ID.
VPI/VCI	Displays the VPI/VCI values.
Admin State	Displays the administrative state of the interface connection.
Oper State	Indicates the status of the ATM interface.
OAM State	Indicates the OAM operational status of ATM connections. ETE — indicates end-to-end connection. AIS — denotes alarm indication signal. RDI — denotes for remote defect indication. LOC — indicates the alarm was due to loss of continuity.
Encap Type	Indicates the encapsulation type.
Owner	Identifies the system entity that owns a specific ATM connection.
AAL Type	Displays ATM Adaptation Layer 5 (AAL5) information.

Table 84 **Output Fields: show port <port-id> atm pvc detail (Continued)**

Label	Description (Continued)
Endpoint Type	Displays the endpoint type.
Cast Type	Indicates the connection topology type.
Type	Indicates the connection type.
Ing. Td Idx	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr. Td Idx	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Last Changed	Indicates the date and time when the interface connection entered its current operational state.
Octets	Displays the number of input and output octets. HEC discarded cells are not included in the input octet numbers.
Cells	Displays the number of input and output cells. HEC discarded cells are not included in the input cell numbers.
Packets	Displays the number of input and output packets. Packets discarded due to HEC or oversize discards are not counted. CRC errors are also in the packet counts show up on the VC level statistics but not on the port level.
Dropped Packets	Displays the number of packets dropped by the ATM SAR device.
CRC-32 Errors	Displays the number of valid AAL-5 SDUs and AAL-5 SDUs with CRC-32 errors received by the AAL-5 VCC.
Reassembly Timeouts	Displays the number of reassembly timeout occurrences.
Over Sized SDUs	Displays the total number of oversized SDU discards.
AIS	Displays the number of AIS cells transmitted and received on this connection for both end to end and segment.
RDI	Displays the number of RDI cells transmitted and received on this connection for both end to end and segment.
Loopback	Displays the number of loopback requests and responses transmitted and received on this connection for both end to end and segment.
CRC-10 Errors	Displays the number of cells discarded on this VPL with CRC 10 errors.

Table 84 Output Fields: show port <port-id> atm pvc detail (Continued)

Label	Description (Continued)
Other	Displays the number of OAM cells that are received but not identified.

Sample Output: show port <port-id> atm pvt detail

```

A:SR1_5>config>service# show port 1/2/2 atm pvt 0.0 detail
=====
ATM PVT
=====
Port Id           : 1/2/2           VPI Range         : 0.0
Admin State       : up              Oper state        : up
Owner             : SAP
Endpoint Type     : PVT             Cast Type         : P2P
Ing. Td Idx       : 1               Egr. Td Idx       : 1
Last Changed      : 04/02/2007 01:59:21
=====
ATM Statistics
=====
                                     Input           Output
-----
Octets                0                   0
Cells                 0                   0
CLP=0 Cells           0                   0
Dropped CLP=0 Cells   0                   0
Dropped Cells (CLP=0+1) 0                   0
Tagged Cells          0
=====
A:SR1_5>config>service#

```

Output Fields: show port <port-id> atm pvt detail

[Table 85](#) describes the output fields for the **show port <port-id> atm pvt detail** command.

Table 85 Output Fields: show port <port-id> atm pvt detail

Label	Description
Port Id	Displays the port ID.
VPI/VCI	Displays the VPI/VCI values.
Admin State	Displays the administrative state of the interface connection.
Oper State	Indicates the status of the ATM interface.
Encap Type	Indicates the encapsulation type.
Owner	Identifies the system entity that owns a specific ATM connection.
Endpoint Type	Displays the endpoint type.

Table 85 Output Fields: show port <port-id> atm pvt detail (Continued)

Label	Description (Continued)
Cast Type	Indicates the connection topology type.
Ingr. Td Idx	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr. Td Idx	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Last Changed	Indicates the date and time when the interface connection entered its current operational state.
Octets	Displays the number of input and output octets. HEC discarded cells are not included in the input octet numbers.
Cells	Displays the number of input and output cells. HEC discarded cells are not included in the input cell numbers.
Dropped CLP	Displays the number of times the CLP1 cells have been dropped. CLP1 cells have lower priority than CLP0 cells and are expected to be discarded first in times of congestion.
Dropped Cells	Displays the number of cells dropped by the ATM SAR device.
Tagged Cells	Displays the number of cells that have been demoted from CLP0 to CLP1.

Sample Output: show port <port-id> [statistics [egress-aggregate]] [detail]

The output displays the aggregate egress queue statistics for ports configured with monitor-aggr-egress-queue-stats which have non-zero counters. This can be shown for a single port, or all ports on an MDA or card. When the detail parameter is added, the output includes those ports with counters that are all zero.

```
*A:PE# show port 2 statistics egress-aggregate
=====
Port 2/1/1 Egress Aggregate Statistics on Slot 2
=====

```

	Forwarded	Dropped	Total
PacketsIn	292	0	292
PacketsOut	0	0	0
OctetsIn	25013	0	25013
OctetsOut	0	0	0

```
=====
Port 2/1/2 Egress Aggregate Statistics on Slot 2
=====

```

	Forwarded	Dropped	Total
PacketsIn	136	0	136
PacketsOut	0	0	0

```

OctetsIn          9320          0          9320
OctetsOut          0          0          0
=====
*A:PE#
*A:PE# show port 2 statistics egress-aggregate detail
=====
Port 2/1/1 Egress Aggregate Statistics on Slot 2
=====
              Forwarded              Dropped              Total
-----
PacketsIn          303              0              303
PacketsOut          0              0              0
OctetsIn           25996            0             25996
OctetsOut           0              0              0
=====
Port 2/1/2 Egress Aggregate Statistics on Slot 2
=====
              Forwarded              Dropped              Total
-----
PacketsIn          140              0              140
PacketsOut          0              0              0
OctetsIn           9598            0             9598
OctetsOut           0              0              0
=====
Port 2/1/3 Egress Aggregate Statistics on Slot 2
=====
              Forwarded              Dropped              Total
-----
PacketsIn           0              0              0
PacketsOut           0              0              0
OctetsIn            0              0              0
OctetsOut           0              0              0
=====
*A:PE#

```

Sample Output: show port statistics

```

=====
Port Statistics on Slot 1
=====
Port      Ingress Packets      Ingress Octets
Id        Egress Packets      Egress Octets
-----
1/1/1          2          1510
              2          1510
1/1/2          2          1510
              2          1510
1/1/3          2          1510
              2          1510
1/1/5          4          1648
              3          1574
1/2/1      18446744073709551608      18446744073709551504
              18446744073709551608      18446744073709551392
1/2/2      18446744073709551608      18446744073709551504
              18446744073709551608      18446744073709551392

```

Output Fields: show port statistics

Table 86 describes the output fields for the **show port statistics** command.

Table 86 Output Fields: show port statistics

Label	Description
Ingress Packets	Displays the ingress packets on a per-port basis.
Ingress Octets	Displays the ingress octets on a per-port basis.
Egress Packets	Displays the egress packets on a per-port basis.
Egress Octets	Displays the egress octets on a per-port basis.

Sample Output: show port <port-id> dot1x

```
A:Dut-C# show port 2/1/11 dot1x
=====
802.1x Port Status
=====
Port control          : force-auth
Port status           : authorized
Authenticator PAE state : force-auth
Backend state         : idle
Reauth enabled        : no           Reauth period          : N/A
Max auth requests     : 2           Transmit period        : 30
Supplicant timeout    : 30          Server timeout         : 30
Quiet period          : 60
Radius-plcy           : N/A
Tunneling              : false
=====
802.1x Session Statistics
=====
authentication method : remote-radius
last session id        : PAC-04258000-6E64D82E
last session time      : 49213d04h
last session username  : N/A
last session term cause : N/A
user tx octets         : 1525626453326
user tx frames         : 329677551
user rx octets         : 1399995911442
user rx frames         : 302536308
=====
Macsec
=====
Admin State           : Up
eapol-destination-address :
Security Zone         : 3
ca-name               : dut_B_C_256_01
=====
```

Output Fields: show port <port-id> dot1x

Table 87 describes the output fields for the following command.

• **show port <port-id> dot1x**

Table 87 Output Fields: show port <port-id> dot1x

Label	Description
Port control	Specifies the 802.1x port control: auto, force-auth, force-unauth.
Port status	Specifies the 802.1x port status.
Authenticator PAE state	Specifies the 802.1x port authenticator PAE state.
Backend state	Specifies the 802.1x port backend state.
Reauth enabled	no — The 802.1x port reauth enabled is not up. yes — The 802.1x port reauth enabled is up.
Reauth period	Specifies the 802.1x port reauthorization period.
Max auth requests	Specifies the 802.1x port maximum authorization requests.
Transmit period	Specifies the 802.1x port transmit period.
Supplicant timeout	Specifies the 802.1x port supplicant timeout.
Server timeout	Specifies the 802.1x port server timeout.
Quiet period	Specifies the 802.1x port quiet period: 1-3600 seconds.
Radius-plcy	Specifies the 802.1x port RADIUS policy name.
Tunneling	true — The 802.1x port tunneling is on. false — The 802.1x port tunneling is not on.
authentication method	Specifies the 802.1x session authentication method.
last session id	Specifies the 802.1x last session ID.
last session time	Specifies the 802.1x last session time.
last session username	Specifies the 802.1x last session username.
last session term cause	Specifies the 802.1x last session term cause.
user tx octets	Specifies the 802.1x session user Tx octets.
user tx frames	Specifies the 802.1x session user Tx frames.
user rx octets	Specifies the 802.1x session user Rx octets.

Table 87 Output Fields: show port <port-id> dot1x (Continued)

Label	Description
user rx frames	Specifies the 802.1x session user Rx frames.
Admin State	Up — The MACsec is administratively up. Down — The MACsec is administratively down. If port <x/y/z> ethernet macsec sub-port 1 is shutdown, the admin state will be down. Otherwise, the admin state is up.
eapol-destination-address	Specifies the destination mac address used in the EAPoL packet for MACsec Key Agreement (MKA) PDUs.
Security Zone	Specifies which security zone this port belongs to.
ca-name	Specifies the CA name assigned to this port.

Sample Output: show port esat-n/n/n

The output displays Ethernet satellite port information.

*A:Dut-A# show port esat-1/1/2

```

=====
Ethernet Interface
=====
Description      : 10/100/Gig Ethernet SFP
Interface        : esat-1/1/2
Link-level       : Ethernet
Admin State      : up
Oper State       : up
Physical Link     : Yes
Single Fiber Mode : No
IfIndex          : 1140918274
Last State Change : 06/24/2016 09:33:16
Last Cleared Time  : N/A
Phys State Chng Cnt: 1

Oper Speed       : 1 Gbps
Config Speed     : 1 Gbps
Oper Duplex      : full
Config Duplex    : full
MTU              : 1578
Min Frame Length : 64 Bytes
Hold time up     : 0 seconds
Hold time down   : 0 seconds
DDM Events       : Enabled

Configured Mode   : access
Dot1Q Ethertype   : 0x8100
PBB Ethertype     : 0x88e7
Ing. Pool % Rate  : 100
Ing. Pool Policy  : n/a
Egr. Pool Policy  : n/a
Net. Egr. Queue Pol: default
Egr. Sched. Pol   : n/a
Monitor Port Sched : Disabled
Monitor Agg Q Stats: Disabled
Auto-negotiate    : true
Oper Phy-tx-clock : not-applicable
Accounting Policy  : None
Acct Plcy Eth Phys : None
Egress Rate       : Default
Load-balance-algo : Default
Access Bandwidth  : Not-Applicable
Access Available BW: 0

Encap Type        : null
QinQ Ethertype    : 0x8100
Egr. Pool % Rate  : 100

MDI/MDX           : MDI
Collect-stats     : Disabled
Collect Eth Phys  : Disabled
Ingress Rate      : Default
LACP Tunnel       : Disabled
Booking Factor    : 100

```

```

Access Booked BW      : 0
Sflow                 : Disabled

Down-when-looped     : Disabled          Keep-alive       : 10
Loop Detected        : False             Retry             : 120
Use Broadcast Addr   : False

Sync. Status Msg.    : Disabled          Rx Quality Level  : N/A
Tx DUS/DNU           : Disabled          Tx Quality Level  : N/A
SSM Code Type        : sdh

Down On Int. Error   : Disabled          DOIE Tx Disable   : N/A

CRC Mon SD Thresh    : Disabled          CRC Mon Window    : 10 seconds
CRC Mon SF Thresh    : Disabled

EFM OAM              : Disabled          EFM OAM Link Mon  : Disabled

Configured Address   : a4:7b:2c:c7:3f:d6
Hardware Address     : a4:7b:2c:c7:3f:d6

```

Transceiver Data

```

Transceiver Status : operational
Transceiver Type   : SFP
Model Number       : 3HE00027CAAA01 ALA IPUIBDKDAA
TX Laser Wavelength: 850 nm              Diag Capable      : yes
Connector Code     : LC                  Vendor OUI         : 00:90:65
Manufacture date    : 2014/11/30         Media             : Ethernet
Serial Number       : NSN1ABK
Part Number         : FTLF8519P3BTL-A5
Optical Compliance : GIGE-SX
Link Length support: 550m for OM2 50u MMF; 300m for OM1 62.5u MMF

```

```

=====
Transceiver Digital Diagnostic Monitoring (DDM), Internally Calibrated
=====

```

	Value	High Alarm	High Warn	Low Warn	Low Alarm
Temperature (C)	+31.7	+110.0	+93.0	-30.0	-40.0
Supply Voltage (V)	3.25	3.60	3.50	3.10	3.00
Tx Bias Current (mA)	6.2	13.0	12.5	2.0	1.0
Tx Output Power (dBm)	-5.08	0.00	-3.00	-9.50	-13.50
Rx Optical Power (avg dBm)	-5.81	0.50	-1.00	-16.99	-21.02

```

=====

```

Traffic Statistics

```

=====

```

	Input	Output
Octets	1585504864	1585783120
Packets	3196585	3197146
Errors	0	0

```

=====

```

Port Statistics

```

=====

```

```

                                     Input          Output
-----
Unicast Packets                     3196583          3197143
Multicast Packets                     0              0
Broadcast Packets                     2              3
Discards                             0              0
Unknown Proto Discards                0
=====

Ethernet-like Medium Statistics
=====

Alignment Errors :                0  Sngl Collisions :                0
FCS Errors       :                0  Mult Collisions :                0
SQE Test Errors  :                0  Late Collisions :                0
CSE              :                0  Excess Collisns :                0
Too long Frames  :                0  Int MAC Tx Errs :                0
Symbol Errors    :                0  Int MAC Rx Errs :                0
=====

```

Sample Output: show port ethernet

The following output is an example of Ethernet information.

A:ALU-1# show port 5/1/4 detail

```

=====
Ethernet Interface
=====
Description      : 10-Gig Ethernet
Interface        : 5/1/4                      Oper Speed      : 10 Gbps
Link-level       : Ethernet                   Config Speed    : N/A
Admin State      : up                        Oper Duplex     : full
Oper State       : down                     Config Duplex    : N/A
Physical Link    : No                       MTU             : 9212
Single Fiber Mode : No                     Min Frame Length : 64 Bytes
IfIndex          : 170000384                Hold time up    : 0 seconds
Last State Change : 04/28/2017 13:09:15      Hold time down  : 0 seconds
Last Cleared Time  : N/A                     DDM Events      : Enabled
Phys State Chng Cnt: 10                     RS-FEC Mode     : None

Configured Mode   : network                  Encap Type      : null
Dot1Q Ethertype   : 0x8100                  QinQ Ethertype   : 0x8100
PBB Ethertype     : 0x88e7
Ing. Pool % Rate  : 100                      Egr. Pool % Rate : 100
Ing. Pool Policy  : n/a
Egr. Pool Policy  : n/a
Net. Egr. Queue Pol: default
Egr. Sched. Pol   : n/a
HS Scheduler Plcy : default
HS Port Pool Plcy : default
Monitor Port Sched : Disabled
Monitor Agg Q Stats: Disabled
Auto-negotiate    : N/A                      MDI/MDX         : N/A
Oper Phy-tx-clock : not-applicable
Accounting Policy  : None                     Collect-stats    : Disabled
Acct Plcy Eth Phys : None                    Collect Eth Phys : Disabled
Egress Rate       : Default                   Ingress Rate     : Default
Load-balance-algo : Default                   LACP Tunnel      : Disabled

```

```

Access Bandwidth   : Not-Applicable           Booking Factor    : 100
Access Available BW: 0
Access Booked BW   : 0
Sflow              : Disabled

Dampening State    : Active                   Current Penalties: 2297
Suppress Threshold : 1500                     Reuse Threshold  : 1000
Max Penalties      : 4000                     Max Suppress Time: 40 seconds
Half Life          : 20 seconds

Down-when-looped   : Disabled                 Keep-alive        : 10
Loop Detected      : False                    Retry             : 120
Use Broadcast Addr : False

Sync. Status Msg.  : Disabled                 Rx Quality Level  : N/A
Tx DUS/DNU         : Disabled                 Tx Quality Level  : N/A
SSM Code Type      : sdh

Down On Int. Error : Disabled                 DOIE Tx Disable   : Disabled

CRC Mon SD Thresh  : Disabled                 CRC Mon Window    : 10 seconds
CRC Mon SF Thresh  : Disabled

Sym Mon SD Thresh  : Disabled                 Sym Mon Window    : 10 seconds
Sym Mon SF Thresh  : Disabled                 Tot Sym Mon Errs  : 0

EFM OAM            : Disabled                 EFM OAM Link Mon  : Disabled
Ignr EFM OAM State : False

Configured Address : 10:e8:78:4d:24:f0
Hardware Address   : 10:e8:78:4d:24:f0
Cfg Alarm          : remote local

```

Transceiver Data

```

Transceiver Status : operational
Transceiver Type    : SFP
Model Number        : 3HE04823AAAA01 ALA IPU3ANKEAA
TX Laser Wavelength: 1310 nm                  Diag Capable      : yes
Connector Code      : LC                      Vendor OUI         : 00:90:65
Manufacture date    : 2016/07/23              Media              : Ethernet
Serial Number       : AW40PC8
Part Number         : FTLX1471D3BCL-A5
Optical Compliance  : 10GBASE-LR
Link Length support: 10km for SMF

```

===== Transceiver Digital Diagnostic Monitoring (DDM), Internally Calibrated =====

	Value	High Alarm	High Warn	Low Warn	Low Alarm
Temperature (C)	+38.1	+78.0	+73.0	-8.0	-13.0
Supply Voltage (V)	3.31	3.70	3.60	3.00	2.90
Tx Bias Current (mA)	40.3	85.0	80.0	20.0	15.0
Tx Output Power (dBm)	-1.29	2.00	1.00	-7.00	-8.00
Rx Optical Power (avg dBm)	-1.38	2.50	2.00	-18.01	-20.00


```

=====
OTU Interface
=====
OTU Status      : Disabled
=====

=====
Traffic Statistics
=====
                                     Input      Output
-----
Octets                0                0
Packets               0                0
Errors                0                0
Utilization (300 seconds)  0.00%        0.00%

=====
Ethernet Statistics
=====
Broadcast Pkts : 0 Drop Events : 0
Multicast Pkts : 0 CRC/Align Errors : 0
Undersize Pkts : 0 Fragments : 0
Oversize Pkts : 0 Jabbers : 0
Collisions : 0
Octets : 0
Packets : 0
Packets of 64 Octets : 0
Packets of 65 to 127 Octets : 0
Packets of 128 to 255 Octets : 0
Packets of 256 to 511 Octets : 0
Packets of 512 to 1023 Octets : 0
Packets of 1024 to 1518 Octets : 0
Packets of 1519 or more Octets : 0

=====

=====
Port Statistics
=====
                                     Input      Output
-----
Unicast Packets                0                0
Multicast Packets              0                0
Broadcast Packets              0                0
Discards                      0                0
Unknown Proto Discards        0

=====

=====
Ethernet-like Medium Statistics
=====
Alignment Errors : 0 Sngl Collisions : 0
FCS Errors : 0 Mult Collisions : 0
SQE Test Errors : 0 Late Collisions : 0
CSE : 0 Excess Collisns : 0
Too long Frames : 0 Int MAC Tx Errs : 0
Symbol Errors : 0 Int MAC Rx Errs : 0
In Pause Frames : 0 Out Pause Frames : 0
=====

```

```

=====
Per Threshold MDA Discard Statistics
=====

```

	Packets	Octets
Threshold 0 Dropped :	0	0
Threshold 1 Dropped :	0	0
Threshold 2 Dropped :	0	0
Threshold 3 Dropped :	0	0
Threshold 4 Dropped :	0	0
Threshold 5 Dropped :	0	0
Threshold 6 Dropped :	0	0
Threshold 7 Dropped :	0	0
Threshold 8 Dropped :	0	0
Threshold 9 Dropped :	0	0
Threshold 10 Dropped :	0	0
Threshold 11 Dropped :	0	0
Threshold 12 Dropped :	0	0
Threshold 13 Dropped :	0	0
Threshold 14 Dropped :	0	0
Threshold 15 Dropped :	0	0

```

=====

Ingress Port Forwarding Engine Drop Reason Statistics
=====

```

IPv4 Header Error	0
IPv4 Invalid Address	0
IPv6 Header Error	0
IPv6 Invalid Address	0
IP Route Blackholed	0
ACL Filter Discards	0
Unicast RPF Check Failed	0
BFD Spoof Check Failed	0
Unicast MAC Destination Address Mismatch	0
Multicast MAC With Unicast Dest IP	0
Unknown MAC Destination Address Discarded in VPLS	0
L2 Service MTU Exceeded	0
Needs ICMP	0

```

=====

Queue Statistics
=====

```

Ingress Queue	Packets	Octets
Ingress Queue 1		
In Profile forwarded :	0	0
In Profile dropped :	0	0
Out Profile forwarded :	0	0
Out Profile dropped :	0	0
Ingress Queue 2		
In Profile forwarded :	0	0
In Profile dropped :	0	0
Out Profile forwarded :	0	0
Out Profile dropped :	0	0
Ingress Queue 3		
In Profile forwarded :	0	0
In Profile dropped :	0	0

```

        Out Profile forwarded :      0      0
        Out Profile dropped   :      0      0
Ingress Queue 4              Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 5              Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 6              Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 7              Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 8              Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 9              Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 10             Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 11             Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 12             Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 13             Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 14             Packets      Octets
    In Profile forwarded    :      0      0
    In Profile dropped      :      0      0
    Out Profile forwarded   :      0      0
    Out Profile dropped     :      0      0
Ingress Queue 15             Packets      Octets
    In Profile forwarded    :      0      0

```

```

        In Profile dropped      :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0
Ingress Queue 16              Packets      Octets
        In Profile forwarded   :      0      0
        In Profile dropped     :      0      0
        Out Profile forwarded   :      0      0
        Out Profile dropped     :      0      0

Egress Queue 1                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 2                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 3                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 4                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 5                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 6                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 7                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
Egress Queue 8                Packets      Octets
        In/Inplus Prof fwded   :      0      0
        In/Inplus Prof dropped:      0      0
        Out/Exc Prof fwded     :      0      0
        Out/Exc Prof dropped   :      0      0
=====

```

Sample Output: show port ethernet efm-oam

The following output is an example of EFM-OAM information.

```

show port 1/1/1 ethernet efm-oam
=====
Ethernet Oam (802.3ah)
=====

```

```

Admin State      : down
Oper State      : disabled
Mode            : active
Pdu Size        : 1518
Config Revision  : 0
Function Support : LB
Transmit Interval : 1000 ms
Multiplier      : 5
Hold Time       : 0
Tunneling       : false
Loop Detected    : false
Grace Tx Enable  : true (inactive)
Grace Vendor OUI : 00:16:4d
Dying Gasp on Reset : true (inactive)
Soft Reset Tx Act : none
Trigger Fault    : none
Vendor OUI       : 00:16:4d (alu)
Vendor Info      : 00:01:00:02
No Peer Information Available
Loopback State   : None
Loopback Ignore Rx : Ignore
Ignore Efm State : false
Link Monitoring  : disabled
Peer RDI Rx
  Critical Event : out-of-service
  Dying Gasp     : out-of-service
  Link Fault     : out-of-service
  Event Notify   : log-only
Local SF Action
  Event Burst    : 1
  Port Action    : out-of-service
  Dying Gasp     : disabled
  Critical Event : disabled
Discovery
  Ad Link Mon Cap : yes
Errored Frame
  Enabled        : no
  Event Notify   : enabled
  SF Threshold   : 1
  SD Threshold   : disabled (0)
  Window         : 10 ds
Errored Symbol Period
  Enabled        : no
  Event Notify   : enabled
  SF Threshold   : 1
  SD Threshold   : disabled (0)
  Window (time)  : 10 ds
  Window (symbols) : 125000000
Errored Frame Period
  Enabled        : no
  Event Notify   : enabled
  SF Threshold   : 1
  SD Threshold   : disabled (0)
  Window         : 1488095 frames
Errored Frame Seconds Summary
  Enabled        : no
  Event Notify   : enabled
  SF Threshold   : 1
  SD Threshold   : disabled (0)
  Window         : 600 ds
=====
Active Failure Ethernet OAM Event Logs
=====
Number of Logs : 0
=====

Ethernet Oam Statistics
=====
                                     Input      Output
-----
Information                          0          0
Loopback Control                     0          0

```

```

Unique Event Notify          0          0
Duplicate Event Notify       0          0
Unsupported Codes            0          0
Frames Lost                  0          0
=====

```

Sample Output: show port ethernet efm-oam event-logs

The following output is an example of EFM-OAM (Link OAM) event logs.

```

show port 1/2/1 ethernet efm-oam event-logs
=====
Active Failure Ethernet OAM Event Logs
=====
Log Index          : 2
Event Time Reference : 10d 03:58:24
Location           : remote
Type               : Dying Gasp
Event Total        : 1
Port Affecting     : yes
-----
Number of Logs : 1
=====

=====
Active Degraded Ethernet OAM Event Logs
=====
Number of Logs : 0
=====

=====
Cleared Failure Ethernet OAM Event Logs
=====
Number of Logs : 0
=====

=====
Cleared Degraded Ethernet OAM Event Logs
=====
Number of Logs : 0
=====

```

Sample Output: show port <port-id> ethernet lldp

```

show port 1/1/1 ethernet lldp
=====
Link Layer Discovery Protocol (LLDP) Port Information
=====

Port 1/1/1 Bridge nearest-bridge
-----
Admin State          : txAndRx          Notifications      : Disabled
Tunnel Nearest Bridge : Disabled
Transmit TLVs        : portDesc sysName sysDesc sysCap
PortID TLV Subtype   : tx-if-name

Management Address Transmit Configuration:
Index 1 (system)     : Enabled          Address             : 10.1.1.31

```

```

Index 2 (IPv6 system) : Disabled      Address      : ::

Port 1/1/1 Bridge nearest-non-tpmr
-----
Admin State           : disabled      Notifications : Disabled
Transmit TLVs         : None
PortID TLV Subtype    : tx-local

Management Address Transmit Configuration:
Index 1 (system)      : Disabled      Address      : 10.1.1.31
Index 2 (IPv6 system) : Disabled      Address      : ::

Port 1/1/1 Bridge nearest-customer
-----
Admin State           : disabled      Notifications : Disabled
Transmit TLVs         : None
PortID TLV Subtype    : tx-local

Management Address Transmit Configuration:
Index 1 (system)      : Disabled      Address      : 10.1.1.31
Index 2 (IPv6 system) : Disabled      Address      : ::

=====

```

Sample Output: show port <port-id> ethernet lldp remote-info

```

show port 1/1/1 ethernet lldp remote-info
=====
Link Layer Discovery Protocol (LLDP) Port Information
=====
Port 1/1/1 Bridge nearest-bridge Remote Peer Information
-----
Remote Peer Index 9 at timestamp 12/08/2014 21:34:30:
Supported Caps       : bridge router
Enabled Caps         : bridge router
Chassis Id Subtype   : 4 (macAddress)
Chassis Id           : D8:1C:FF:00:00:00
PortId Subtype       : 5 (interfaceName)
Port Id              : 31:2F:32:2F:32
                     "1/2/2"
Port Description     : n/a
System Name          : cses-V28
System Description    : TiMOS-B-0.0.I4269 both/i386 Nokia 7750 SR Copyright
                       (c) 2000-2016 Nokia.
                       All rights reserved. All use subject to applicable
                       license agreements.
                       Built on Wed Dec 3 19:14:27 PST 2014 by builder in /
                       rel0.0/I4269/panos/main

Port 1/1/1 Bridge nearest-non-tpmr Remote Peer Information
-----
No remote peers found

Port 1/1/1 Bridge nearest-customer Remote Peer Information
-----
No remote peers found

=====

```

Sample Output: show port <port-id> ethernet lldp remote-info detail

```

show port 1/1/1 ethernet lldp remote-info detail
=====
Link Layer Discovery Protocol (LLDP) Port Information
=====
Port 1/1/1 Bridge nearest-bridge Remote Peer Information
-----
Remote Peer Index 9 at timestamp 12/08/2014 21:34:30:
Supported Caps      : bridge router
Enabled Caps       : bridge router
Chassis Id Subtype  : 4 (macAddress)
Chassis Id         : D8:1C:FF:00:00:00
PortId Subtype     : 5 (interfaceName)
Port Id            : 31:2F:32:2F:32
                   "1/2/2"
Port Description    : n/a
System Name        : cses-V28
System Description  : TiMOS-B-0.0.I4269 both/i386 Nokia 7750 SR Copyright
                   (c) 2000-2016 Nokia.
                   All rights reserved. All use subject to applicable
                   license agreements.
                   Built on Wed Dec 3 19:14:27 PST 2014 by builder in /
                   rel0.0/I4269/panos/main

Remote Peer Index 9 management addresses at time 12/08/2014 21:34:30:
Address SubType     : 1 (IPv4)
Address             : 10.1.1.28
Address If SubType  : 2           Address If Id           : 1
Address OID         : .1.3.6.1.4.1.6527.1.3.3

Port 1/1/1 Bridge nearest-non-tpmr Remote Peer Information
-----
No remote peers found

Port 1/1/1 Bridge nearest-customer Remote Peer Information
-----
No remote peers found

=====

```

Sample Output: show port <port-id> ethernet lldp detail

```

show port 1/1/1 ethernet lldp detail
=====
Link Layer Discovery Protocol (LLDP) Port Information
=====
Port 1/1/1 Bridge nearest-bridge
-----
Admin State          : txAndRx          Notifications      : Disabled
Tunnel Nearest Bridge : Disabled
Transmit TLVs        : portDesc sysName sysDesc sysCap
PortID TLV Subtype   : tx-if-name

Management Address Transmit Configuration:
Index 1 (system)     : Enabled          Address             : 10.1.1.31

```



```

Index 2 (IPv6 system) : Disabled      Address      : ::

Port LLDP Stats:
Tx Frames              : 11749        Tx Length Err Frames : 0
Rx Frames              : 70399        Rx Frame Discard     : 0
Rx Frame Errors        : 0            Rx TLV Discard       : 0
Rx TLV Unknown         : 0            Rx Ageouts           : 3

Port 1/1/1 Bridge nearest-non-tpmr
-----
Admin State            : disabled      Notifications     : Disabled
Transmit TLVs          : None
PortID TLV Subtype     : tx-local

Management Address Transmit Configuration:
Index 1 (system)       : Disabled      Address           : 10.1.1.31
Index 2 (IPv6 system) : Disabled      Address           : ::

Port LLDP Stats:
Tx Frames              : 0            Tx Length Err Frames : 0
Rx Frames              : 0            Rx Frame Discard     : 0
Rx Frame Errors        : 0            Rx TLV Discard       : 0
Rx TLV Unknown         : 0            Rx Ageouts           : 0

Port 1/1/1 Bridge nearest-customer
-----
Admin State            : disabled      Notifications     : Disabled
Transmit TLVs          : None
PortID TLV Subtype     : tx-local

Management Address Transmit Configuration:
Index 1 (system)       : Disabled      Address           : 10.1.1.31
Index 2 (IPv6 system) : Disabled      Address           : ::

Port LLDP Stats:
Tx Frames              : 0            Tx Length Err Frames : 0
Rx Frames              : 0            Rx Frame Discard     : 0
Rx Frame Errors        : 0            Rx TLV Discard       : 0
Rx TLV Unknown         : 0            Rx Ageouts           : 0

=====

```

Sample Output: show port <port-id> macsec sub-port <sub-port-id>

```

A:Dut-C# show port 2/1/11 macsec sub-port 1
=====
Port 2/1/11:1 MACsec
=====
Admin State            : Up
eapol-destination-address :
Security Zone          : 3
CA Name                : dut_B_C_256_01
=====
Connectivity Association "dut_B_C_256_01"
=====

```

```

Admin State      : Up
Description      : (Not Specified)
Replay Protection : Enabled
Replay Window Size : 0
Macsec Encrypt   : Enabled
Clear Tag Mode    : none
Cipher Suite      : gcm-aes-256
Encryption Offset : 0
Assigned ports    : 2/1/11
-----
Static Cak
-----
MKA Key Server Priority : 16
Active Pre-Shared-Key Index : 1
Active Pre-Shared-Key CKN : 11223344556677889900aabbccddeeff112233445566778*
=====
Dot1x macsec
=====
Subport <x>
Encap-match: encap
Service-id port-id: 1 1/1/c1/1:1
Interface-name port-id: na
eapol-destination-address : dest mac
macsec-enabled : yes
Security zone <zone>
ca-name <ca-name>
=====
* indicates that the corresponding row element may have been truncated.

```

Output Fields: show port <port-id> macsec subport-id <sub-port-id>

[Table 88](#) describes the output fields for the **show port <port-id> macsec subport-id <sub-port-id>** command.

Table 88 **Output Fields: show port <port-id> macsec sub-port <sub-port-id>**

Label	Description
Admin State	Up — The CA is administratively up. Down — The CA is administratively down. If port <x/y/z> ethernet>macsec sub-port 1 is shutdown, the admin state will be down. Otherwise, the admin state is up.
eapol-destination-address	Specifies the destination mac address used in the EAPoL packet for MACsec Key Agreement (MKA) PDUs.
Security Zone	Specifies which security zone this port belongs to. Security zones provide the scalability restrictions in terms of MACsec peers per zone.
CA name	Specifies the CA name.
Description	Specifies a user description for this CA.

Table 88 **Output Fields: show port <port-id> macsec sub-port <sub-port-id> (Continued)**

Label	Description
Replay Protection	Enabled — Replay Protection is enabled. Disabled — Replay Protection is disabled. If replay protection is enabled for this CA, the out of replay-window packets will be discarded. Otherwise, the Replay Window Size value is ignored. The default value is disabled, so that when adding MACsec to an interface, it will not affect traffic immediately.
Replay Window Size	Specifies the size, in packets, of the replay window. Each packet is assigned a unique packet number. Replay protection enforces strict ordering of the packets and protects against replay attacks. In networks where out-of-order packets are expected, the window size can be set up to 4294967295.
Macsec Encrypt	Enabled — MACsec encryption is enabled on the traffic. All fields behind source or destination MAC addresses are encrypted on each packet. Disabled — MACsec encryption is disabled for the traffic. The packet fields are sent in clear text but data integrity is checked on each packet, and the MACsec overhead (header and ICV) is still added to each packet.
Clear Tag Mode	In the case of VLAN-tagged traffic, if the traffic is crossing a network, one or two VLAN tags can be transmitted in clear text so that this traffic can receive preferential treatment over the network based on 802.1 q-tag or p-bits. Values: single-tag, dual-tag
Cipher Suite	Specifies the cipher suite used for encrypting the SAK: gcm-aes-128, gcm-aes-256, gcm-aes-xpn-128, gcm-aes-xpn-256.
Encryption Offset	Specifies the encryption offset configured on this node: 0, 30, 50. In the case of IP traffic, the IP header can be transmitted in clear text so that the traffic is routed accordingly when crossing the network. A value of 30 will be used for IPv4 and 50 for IPv6.
Assigned ports	Specifies all ports that contain this CA.
MKA Key Server Priority	Specifies the MKA key server priority: 0 to 255 (default 16). A priority of 0 means more it is more likely to become the MKA server. If multiple participants have the same priority, the MAC address becomes the differentiator for electing the MKA server.
Active Pre-Shared Key Index	Specifies the active pre-shared key index: 1-2 (default 1).

Table 88 **Output Fields: show port <port-id> macsec sub-port <sub-port-id> (Continued)**

Label	Description
Active Pre-Shared Key CKN	Specifies the active PSK CAK name.
Subport	Specifies the MACsec sub-port used. Multiple MACsec instances can be configured under one port, matching traffic based on the encap-match parameter.
Encap-match	Specifies which types of traffic MACsec is enabled for.
Service-id port-id	Specifies the service IDs affected by the encap-match configured under this sub-port. One or more services can be affected by the match (by using * or by targeting ports with traffic aggregations, for example).
Interface-name port-id	Specifies the interface name of the port.
eapol-destination-address	Specifies the destination MAC address used in the EAPoL packet for MACsec Key Agreement (MKA) PDUs.
macsec-enabled	Specifies if MACSec is enabled.
Security zone	Specifies which security zone this port belongs to. Security zones provide scalability restrictions in terms of MACsec peers per zone.
ca-name	Specifies the CA name assigned to this port.

Sample Output: show port <port-id> macsec sub-port <sub-port-id> detail

```

*A:Dut-C# show port 2/1/11 macsec sub-port 1 statistics
=====
MACsec Statistics
=====
-----
txSecyStats
-----
Untagged Packets           : 0
Too Long Packets          : 0
-----
txSCSecyStats
-----
TxSc SCI PortNum : 1
Protected Packets : 0
Encrypted Packets : 0
Protected Octets  : 0
Encrypted Octets  : 0
-----
txSASecyStats
-----
AN : 0

```

```

        Protected SA Packets          : 0
        Encrypted SA Packets          : 0
    AN : 1
        Protected SA Packets          : 0
        Encrypted SA Packets          : 0
    AN : 2
        Protected SA Packets          : 0
        Encrypted SA Packets          : 0
    AN : 3
        Protected SA Packets          : 0
        Encrypted SA Packets          : 0
-----
rxSecyStats
-----
        No Tag Packets                : 4681
        Bad Tag Packets               : 0
        No SCI Packets                : 0
        Overrun Packets               : 0
-----
rxSCSecyStats
-----
SCI : a47b2ce1110d0001
        No Using SA Packets          : 0
        Late Packets                  : 0
        Not Valid Packets             : 0
        Delayed Packets                : 0
        Unchecked Packets              : 0
        OK Packets                    : 0
        Validated Octets               : 0
        Decrypted Octets               : 0
-----
rxSASecyStats
-----
SCI : a47b2ce1110d0001
    AN : 2
        No Using SA Packets          : 0
        Not Valid Packets             : 0
        OK Packets                    : 0
    AN : 3
        No Using SA Packets          : 0
        Not Valid Packets             : 0
        OK Packets                    : 0
-----
rxSCSecyStats
-----
SCI : a47b2ce1128e0001
        No Using SA Packets          : 0
        Late Packets                  : 0
        Not Valid Packets             : 0
        Delayed Packets                : 0
        Unchecked Packets              : 0
        OK Packets                    : 0
        Validated Octets               : 0
        Decrypted Octets               : 0
-----
rxSASecyStats
-----
SCI : a47b2ce1128e0001
    AN : 0

```

```

        No Using SA Packets      : 0
        Not Valid Packets       : 0
        OK Packets               : 0
AN : 1
        No Using SA Packets      : 0
        Not Valid Packets       : 0
        OK Packets               : 0
AN : 2
        No Using SA Packets      : 0
        Not Valid Packets       : 0
        OK Packets               : 0
AN : 3
        No Using SA Packets      : 0
        Not Valid Packets       : 0
        OK Packets               : 0
=====

```

Sample Output: show port <port-id> macsec sub-port <sub-port-id> statistics

```

A:Dut-C# show port 2/1/11 macsec statistics
=====
MACsec Statistics
=====
-----
txSecyStats
-----
        Untagged Packets      : 0
        Too Long Packets     : 0
-----
rxSecyStats
-----
        No Tag Packets        : 758
        Bad Tag Packets       : 0
        No SCI Packets        : 0
        Overrun Packets       : 0
-----
txSCSecyStats
-----
        Protected Packets     : 0
        Encrypted Packets     : 355170025
        Protected Octets      : 0
        Encrypted Octets      : 1637917300250
-----
rxSCSecyStats
-----
SCI : 00:00:a4:7b:2c:e1
        No Using SA Packets   : 0
        Late Packets          : 0
        Not Valid Packets     : 0
        Delayed Packets       : 0
        Unchecked Packets     : 0
        OK Packets            : 325904694
        Validated Octets      : 0
        Decrypted Octets      : 1502922287478
-----
txSASecyStats
-----
AN : 0

```

```

Protected SA Packets      : 0
Encrypted SA Packets     : 355170201
-----
rxSASecyStats
-----
SCI : 00:00:a4:7b:2c:e1
  AN : 0
    No Using SA Packets   : 0
    Not Valid Packets    : 0
    OK Packets           : 325904947
=====

```

Output Fields: show port <port-id> macsec sub-port <sub-port-id> statistics

Table 89 describes the output fields for the following command.

- show port <port-id> macsec sub-port <sub-port-id> statistics

Table 89 **Output Fields: show port <port-id> macsec sub-port <sub-port-id> statistics**

Label	Description
Untagged Packets	Indicates the number of transmitted packets without the MAC security tag (SecTAG) when the value of <code>oftmnxMacsecConnAssocReplayProtect</code> for the configured CA is set to 'false'.
Too Long Packets	Indicates the number of transmitted packets discarded because the packet length is greater than the Maximum Transmission Unit (MTU) of the Ethernet physical interface.
No Tag Packets	Indicates the number of received packets discarded without the MAC security tag (SecTAG).
Bad Tag Packets	Indicates the number of received packets discarded with an invalid SecTAG or a zero value Packet Number (PN) or an invalid Integrity Check Value (ICV).
No SCI Packets	Indicates the number of received packets discarded with unknown SCI information when the C bit in the SecTAG is set.
Overrun Packets	Indicates the number of packets discarded because the number of received packets exceeded the cryptographic performance capabilities.
Protected Packets	Indicates the number of packets that are integrity protected but not encrypted for this transmitting SA.
Encrypted Packets	Indicates the number of packets that are integrity protected and encrypted for this transmitting SA.
Protected Octets	Indicates the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.

Table 89 **Output Fields: show port <port-id> macsec sub-port <sub-port-id> statistics (Continued)**

Label	Description
Encrypted Octets	Indicates the number of plain text octets that are integrity protected and encrypted on the transmitting SC.
SCI	Indicates the Secure Channel Identifier (SCI).
No Using SA Packets	Indicates the number of received packets that have been discarded on this SA which is not currently in use.
Late Packets	Indicates the number of received packets that have been discarded due to replay window protection on this SC.
Not Valid Packets	Indicates the number of not valid packets that have been discarded on this active SA.
Delayed Packets	Indicates the number of received packets with the condition a PN lower than the lower bound of the replay protection on this SC.
Unchecked Packets	Indicates the number of packets that have failed the integrity check on this SC.
OK Packets	Indicates the number of validated packets on this SA.
Validated Octets	Indicates the number of octets of plain text recovered from received packets that were integrity protected but not encrypted.
Decrypted Octets	Indicates the number of octets of plain text recovered from received packets that were integrity protected and encrypted.
AN	Indicates the AN for identifying the receiving SA.
Protected SA Packets	Indicates the number of packets that are integrity protected but not encrypted for this SA.
Encrypted SA Packets	Indicates the number of packets that are integrity protected and encrypted for this SA.

port-tree

Syntax	port-tree <i>port-id</i>
Context	show
Description	This command displays the tree for SONET/SDH or TDM ports/channels WAN PHY mode (xgig wan) Ethernet ports.
Parameters	<i>port-id</i> — Specifies the physical port ID.
Values	slot/mda/port [.channel]

Table 90 Port ID Values

Product	Slot	MDA	CMA	Port	Values
7750 SR-12	1 to 10	1, 2		1 to 60 (depending on the MDA type)	—
7750 SR-c12	1	1, 3, 5, 7, 9, 11	1 to 12		—
7750 SR-c4	1	1, 3	1 to 4		—
7750 SR-7	1 to 5	1, 2	—		—
7950 XRS	1 to 20	—	—		—
7450 ESS-7	—		—		1 to 4
7450 ESS-12	—		—		1 to 10
Channelized MDAs					
CHOC12-SFP	—				slot/mda/port. [1 to 4] . [1 to 3] . [1 to 28] . [..24] For example, 7/2/1.1.1.28.24
CHOC3-SFP	—				slot/mda/port. [1 to 3] . [1 to 28] . [..24] For example, 7/2/1.1.28.24
DS3	—				slot/mda/port. [1 to 28] . [..24] For example, 7/1/1.1.1

Output The following output is an example of port information, and [Table 91](#) describes the output fields.

Sample Output: show port-tree <port-id>

```
A:ALA-48>config# show port-tree 7/1/1
  ifIndex  type, sonet-sdh-index (* = provisioned)
=====
119570432  Port, N/A *
656441345  DS3, none *
656441405  DS1, 1 *
656441430  DS1, 2
656441455  DS1, 3
656441480  DS1, 4
656441505  DS1, 5
656441530  DS1, 6
656441555  DS1, 7
656441580  DS1, 8
656441605  DS1, 9
656441630  DS1, 10
656441655  DS1, 11
656441680  DS1, 12
```

```

656441705      DS1, 13
656441730      DS1, 14
656441755      DS1, 15
656441780      DS1, 16
656441805      DS1, 17
656441830      DS1, 18
656441855      DS1, 19
656441880      DS1, 20
656441905      DS1, 21
656441930      DS1, 22
656441980      DS1, 24
656442005      DS1, 25
656442030      DS1, 26
656442055      DS1, 27
656442080      DS1, 28

```

```
A:ALA-48>config#
```

Table 91 **Output Fields: show port-tree <port-id>**

Label	Description
lIndex	Displays the interface index number which reflects its initialization sequence.
type	Specifies the type.
sonet-sdh-index	Specifies the sonet-sdh-index.
*	When an asterisk (*) is displayed after the sonet-sdh-index, the port/channel is provisioned.

redundancy

Syntax **redundancy**

Context show

Description This command enables the context to show multi-chassis redundancy information.

multi-chassis

Syntax **multi-chassis**

Context show>redundancy

Description This command displays multi-chassis redundancy information.

Output See the following sections for output samples:

- [Sample Output: show redundancy multi-chassis mc-lag peer <ip-address>](#)

- [Sample Output: show redundancy multi-chassis mc-lag statistics](#)
- [Sample Output: show redundancy multi-chassis mc-lag peer <ip-address> lag 2 statistics](#)
- [Sample Output: show redundancy multi-chassis mc-lag peer 10.10.10.102 statistics](#)
- [Sample Output: show redundancy multi-chassis sync](#)
- [Sample Output: show redundancy multi-chassis sync peer <ip-address>](#)
- [Sample Output: show redundancy multi-chassis sync peer <ip-address> detail](#)
- [Sample Output: show redundancy multi-chassis sync statistics](#)
- [Sample Output: show redundancy multi-chassis sync peer <ip-address> statistics](#)

Sample Output: show redundancy multi-chassis mc-lag peer <ip-address>

```
*A:Dut-C# show redundancy multi-chassis mc-lag peer 10.10.10.1
=====
Multi-Chassis MC-Lag Peer 10.10.10.1
=====
Last State chg: 09/24/2007 07:58:03
Admin State: Up      Oper State   : Up
KeepAlive: 10 deci-seconds      Hold On Ngbr Failure : 3
-----
Lag Id LACP Key Remote Lag Id System Id  Sys Prio Last State Changed
-----
1      326661      00:00:00:33:33:33  32888  09/24/2007 07:56:35
-----
Number of LAGs : 1
=====
*A:Dut-C#
```

Sample Output: show redundancy multi-chassis mc-lag statistics

```
A:pc1# show redundancy multi-chassis mc-lag statistics
=====
Multi-Chassis Statistics
=====
Packets Rx                : 129816
Packets Rx Keepalive      : 129798
Packets Rx Config         : 3
Packets Rx Peer Config    : 5
Packets Rx State          : 10
Packets Dropped KeepaliveTask : 0
Packets Dropped Packet Too Short : 0
Packets Dropped Verify Failed : 0
Packets Dropped Tlv Invalid Size : 0
Packets Dropped Out of Seq : 0
Packets Dropped Unknown Tlv : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped MD5       : 0
Packets Dropped Unknown Peer : 0
Packets Tx                : 77918
Packets Tx Keepalive      : 77879
Packets Tx Config         : 6
Packets Tx Peer Config    : 26
Packets Tx State          : 7
```

```

Packets Tx Failed          : 0
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis mc-lag peer <ip-address> lag 2 statistics

```

A:pc1# show redundancy multi-chassis mc-lag peer 10.10.10.102 lag 2 statistics
=====
Multi-Chassis Statistics, Peer 10.10.10.102 Lag 2
=====
Packets Rx Config          : 1
Packets Rx State           : 4
Packets Tx Config          : 2
Packets Tx State           : 3
Packets Tx Failed          : 0
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis mc-lag peer 10.10.10.102 statistics

```

A:pc1#show redundancy multi-chassis mc-lag peer 10.10.10.102 statistics
=====
Multi-Chassis Statistics, Peer 10.10.10.102
=====
Packets Rx                  : 129918
Packets Rx Keepalive        : 129900
Packets Rx Config           : 3
Packets Rx Peer Config      : 5
Packets Rx State            : 10
Packets Dropped State Disabled : 0
Packets Dropped Packets Too Short : 0
Packets Dropped Tlv Invalid Size : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped Out of Seq   : 0
Packets Dropped Unknown Tlv  : 0
Packets Dropped MD5          : 0
Packets Tx                  : 77979
Packets Tx Keepalive        : 77940
Packets Tx Peer Config      : 26
Packets Tx Failed           : 0
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis sync

```

A:pc1# show redundancy multi-chassis sync
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 10.10.10.102
Description          : CO1
Authentication       : Enabled
Source IP Address    : 10.10.10.101

```

```

Admin State           : Enabled
-----
Sync-status
-----
Client Applications   :
Sync Admin State      : Up
Sync Oper State       : Up
DB Sync State         : inSync
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
=====
Peer
-----
Peer IP Address       : 10.10.20.1
Authentication        : Disabled
Source IP Address     : 0.0.0.0
Admin State           : Disabled
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis sync peer <ip-address>

```

pc1# show redundancy multi-chassis sync peer 10.10.10.102
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address       : 10.10.10.102
Description           : CO1
Authentication        : Enabled
Source IP Address     : 10.10.10.101
Admin State           : Enabled
-----
Sync-status
-----
Client Applications   :
Sync Admin State      : Up
Sync Oper State       : Up
DB Sync State         : inSync
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
Rem Num Entries       : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries     : 0
=====
MCS Application Stats
=====
Application           : igmp
Num Entries           : 0
Lcl Deleted Entries   : 0
Alarm Entries         : 0
-----

```

```

Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
-----
Application          : igmpSnooping
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
-----
Application          : subMgmt
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
-----
Application          : srrp
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
-----
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries    : 0
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis sync peer <ip-address> detail

```

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 detail
=====
Multi-chassis Peer Table
=====
Peer
-----
Peer IP Address      : 10.10.10.102
Description          : CO1
Authentication       : Enabled
Source IP Address    : 10.10.10.101
Admin State          : Enabled
-----
Sync-status
-----
Client Applications  :
Sync Admin State     : Up
Sync Oper State      : Up
DB Sync State        : inSync
Num Entries          : 0
Lcl Deleted Entries  : 0
Alarm Entries        : 0
Rem Num Entries      : 0
Rem Lcl Deleted Entries : 0

```

```

Rem Alarm Entries      : 0
=====
MCS Application Stats
=====
Application            : igmp
Num Entries            : 0
Lcl Deleted Entries    : 0
Alarm Entries          : 0
-----
Rem Num Entries        : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries      : 0
-----
Application            : igmpSnooping
Num Entries            : 0
Lcl Deleted Entries    : 0
Alarm Entries          : 0
-----
Rem Num Entries        : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries      : 0
-----
Application            : subMgmt
Num Entries            : 0
Lcl Deleted Entries    : 0
Alarm Entries          : 0
-----
Rem Num Entries        : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries      : 0
-----
Application            : srrp
Num Entries            : 0
Lcl Deleted Entries    : 0
Alarm Entries          : 0
-----
Rem Num Entries        : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries      : 0
=====
Ports synced on peer 10.10.10.102
=====
Port/Encap              Tag
-----
1/1/1
  1-2                    r1
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis sync statistics

```

A:pc1# show redundancy multi-chassis sync statistics
=====
Multi-chassis Peer Sync Stats
=====
Peer IP Address        : 10.10.10.102
Packets Tx Total       : 511
Packets Tx Hello       : 510

```

```

Packets Tx Data      : 0
Packets Tx Other     : 1
Packets Tx Error     : 0
Packets Rx Total     : 511
Packets Rx Hello     : 510
Packets Rx Data      : 0
Packets Rx Other     : 1
Packets Rx Error     : 0
Packets Rx Header Err : 0
Packets Rx Body Err  : 0
Packets Rx Seq Num Err : 0
=====
Peer IP Address      : 10.10.20.1
Packets Tx Total     : 0
Packets Tx Hello     : 0
Packets Tx Data      : 0
Packets Tx Other     : 0
Packets Tx Error     : 0
Packets Rx Total     : 0
Packets Rx Hello     : 0
Packets Rx Data      : 0
Packets Rx Other     : 0
Packets Rx Error     : 0
Packets Rx Header Err : 0
Packets Rx Body Err  : 0
Packets Rx Seq Num Err : 0
=====
A:pc1#

```

Sample Output: show redundancy multi-chassis sync peer <ip-address> statistics

```

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 statistics
=====
Multi-chassis Peer Sync Stats
=====
Peer IP Address      : 10.10.10.102
Packets Tx Total     : 554
Packets Tx Hello     : 553
Packets Tx Data      : 0
Packets Tx Other     : 1
Packets Tx Error     : 0
Packets Rx Total     : 554
Packets Rx Hello     : 553
Packets Rx Data      : 0
Packets Rx Other     : 1
Packets Rx Error     : 0
Packets Rx Header Err : 0
Packets Rx Body Err  : 0
Packets Rx Seq Num Err : 0
=====
A:pc1#

```

all

Syntax all

- Context** show>redundancy>multi-chassis
- Description** This command displays all multi-chassis information.
- Output** The following output is an example of multi-chassis all command information.

Sample Output: show redundancy multi-chassis all

```
A:pc1# show redundancy multi-chassis all
=====
Multi-Chassis Peers
=====
Peer IP          Src IP          Auth          Peer Admin
MCS Admin      MCS Oper      MCS State      MC-LAG Admin  MC-LAG Oper
-----
10.10.10.102    10.10.10.101  hash          Enabled
Enabled          Enabled      inSync        Enabled        Enabled
10.10.20.1      0.0.0.0       None          Disabled
--              --            --            Disabled        Disabled
=====
A:pc1#
```

mc-lag

- Syntax** **mac-lag peer** *ip-address* [**lag** *lag-id*]
mac-lag [**peer** *ip-address* [**lag** *lag-id*]] **statistics**
- Context** show>redundancy>multi-chassis
- Description** This command displays multi-chassis LAG information.
- Parameters** *ip-address* — Specifies the IP address.
- Values** ipv4-address: a.b.c.d
ipv6-address:
• x:x:x:x:x:x:x (eight 16-bit pieces)
• x:x:x:x:x:d.d.d.d
• x: [0 to FFFF] H
• d: [0 to 255] D
- lag-id* — Specifies the LAG ID.
- Values** 1 to 800
- statistics** — Displays statistics.
- Output** The following output is an example of multi-chassis LAG information.

Sample Output: show redundancy multi-chassis mc-lag peer <ip-address>

```
*A:Dut-B# show redundancy multi-chassis mc-lag peer 10.20.1.2
=====
```

```

Multi-Chassis MC-Lag Peer 10.20.1.2
=====
Last State chg : 05/17/2009 19:31:58
Admin State : Up Oper State : Up
KeepAlive : 5 deci-seconds Hold On Ngbr Failure : 2
-----
Lag Id Lacp Remote Source Oper System Id Sys Last State Changed
Key Lag Id MacLSB MacLSB Prio
-----
1 40000 1 Lacp 9c:40 00:02:80:01:00:01 100 05/17/2009 19:31:56

*A:Dut-B# /tools dump redundancy src-bmac-lsb
Src-bmac-lsb: 1025 (04-01) User: B-Vpls - 1 service(s)
Services affected:
B-Vpls: 1
B-Vpls: 2

```

mc-ring

Syntax	mc-ring peer <i>ip-address</i> statistics mc-ring peer <i>ip-address</i> [ring <i>sync-tag</i> [detail statistics]] mc-ring peer <i>ip-address</i> ring <i>sync-tag</i> ring-node [<i>ring-node-name</i> [detail statistics]] mc-ring global-statistics
Context	show>redundancy>multi-chassis
Description	This command displays multi-chassis ring information.
Parameters	<i>ip-address</i> — Specifies the address of the multi-chassis peer to display. Values ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF] H • d: [0 to 255] D <i>sync-tag</i> — Specifies a synchronization tag to be displayed that was used while synchronizing this port with the multi-chassis peer up to 32 characters. <i>ring-node-name</i> — Specifies a ring-node name up to 32 characters. global-statistics — Displays global statistics for the multi-chassis ring. detail — Displays detailed peer information for the multi-chassis ring.
Output	See the following sections for output samples: • Sample Output: show redundancy multi-chassis mc-ring peer <ip-address> ring <sync-tag> detail • Output Fields: show redundancy multi-chassis mc-ring peer <ip-address> ring • Sample Output: show redundancy multi-chassis mc-ring peer <ip-address> statistics

- [Output Fields: show redundancy multi-chassis mc-ring peer <ip-address> statistics](#)
- [Output Fields: show redundancy multi-chassis mc-ring global-statistics](#)
- [Output Fields: show redundancy multi-chassis mc-ring ring-node](#)
- [Sample Output: show redundancy multi-chassis mc-ring global-statistics](#)

Sample Output: show redundancy multi-chassis mc-ring peer <ip-address> ring <sync-tag> detail

```
*A:ALA-48# show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 detail
=====
Multi-Chassis MC-Ring Detailed Information
=====
Peer           : 10.0.0.2
Sync Tag       : ring11
Port ID        : 1/1/3
Admin State    : inService
Oper State     : connected
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:24
Last Debounce  : 02/15/2008 09:28:42
Debounce Period: 0d 00:00:00
Failure Reason : None
-----
In Band Control Path
-----
Service ID      : 10
Interface Name  : to_an1
Oper State     : connected
Dest IP        : 10.10.0.2
Src IP         : 10.10.0.1
-----
VLAN Map B Path Provisioned
-----
range 13-13
range 17-17
-----
VLAN Map Excluded Path Provisioned
-----
range 18-18
-----
VLAN Map B Path Operational
-----
range 13-13
range 17-17
-----
VLAN Map Excluded Path Operational
-----
range 18-18
=====
*A:ALA-48#

*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.168.10.104
=====
MC Ring entries
=====
```

```

Sync Tag                      Oper State      Failure Reason
-----
No. of MC Ring entries: 0
=====
*A:ALA-48#

*A:ALA-48# show redundancy multi-chassis mc-ring peer 10.0.0.2
=====
MC Ring entries
=====
Sync Tag                      Oper State      Failure Reason
-----
ring11                       connected       None
ring12                       shutdown        None
-----
No. of MC Ring entries: 4
=====
*A:ALA-48#

*A:ALA-48# show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-
node an1 detail
=====
Multi-Chassis MC-Ring Node Detailed Information
=====
Peer           : 10.0.0.2
Sync Tag       : ring11
Node Name      : an1
Oper State Loc : connected
Oper State Rem : notTested
In Use         : True
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:25
Failure Reason : None
-----
Ring Node Connectivity Verification
-----
Admin State     : inService
Service ID      : 11
VLAN Tag        : 11
Dest IP         : 10.11.3.1
Src IP          : None
Interval        : 1 minutes
Src MAC         : None
=====
*A:ALA-48#

*A:ALA-48# show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node
=====
MC Ring Node entries
=====
Name              Loc Oper St.      Failure Reason
  In Use          Rem Oper St.
-----
an1               connected         None
  Yes             notTested
an2               connected         None

```

```

Yes                                notTested
-----
No. of MC Ring Node entries: 2
=====
*A:ALA-48#

```

Output Fields: `show redundancy multi-chassis mc-ring peer <ip-address> ring`

[Table 92](#) describes the output fields for the `show redundancy multi-chassis mc-ring peer <ip-address> ring` command.

Table 92 **Output Fields: show redundancy multi-chassis mc-ring peer <ip-address> ring**

Label	Description
Sync Tag	Displays the synchronization tag that was used while synchronizing this port with the multi-chassis peer.
Oper State	noPeer — The peer has no corresponding ring configured. connected — The in-band control connection with the peer is operational. broken — The in-band control connection with the peer has timed out. conflict — The in-band control connection with the peer has timed out but the physical connection is still OK; the failure of the in-band signaling connection is caused by a misconfiguration. For example, a conflict between the configuration of this system and its peer, or a misconfiguration on one of the ring access node systems. testingRing — The in-band control connection with the peer is being set up. Waiting for result. waitingForPeer — Verifying if this ring is configured on the peer. configErr — The ring is administratively up, but a configuration error prevents it from operating properly. halfBroken — The in-band control connection indicates that the ring is broken in one direction (towards the peer). localBroken — The in-band control connection with the peer is known to be broken due to local failure or local administrative action. shutdown — The ring is shutdown.
Failure Reason	Displays the failure reason.
Last Debounce	Displays the last time that the debounce mechanism (protecting the router from overload situations in case of a flapping ring) was activated.
Debounce Period	Displays the duration that the debounce mechanism was in action since the “Last Debounce”.

Sample Output: `show redundancy multi-chassis mc-ring peer <ip-address> statistics`

```

*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.168.10.104 statistics
=====
MC Ring statistics for peer 192.168.10.104

```

```

=====
Message                                     Received      Transmitted
-----
MCS ID Request                            0             0
MCS ID Response                           0             0
Ring Exists Request                       0             0
Ring Exists Response                      0             0
Keepalive                                0             0
-----
Total                                    0             0
=====
*A:ALA-48>show>redundancy>multi-chassis#

```

Output Fields: show redundancy multi-chassis mc-ring peer <ip-address> statistics

Table 93 describes the output fields for the **show redundancy multi-chassis mc-ring peer <ip-address> statistics** command.

Table 93 **Output Fields: show redundancy multi-chassis mc-ring peer <ip-address> statistics**

Label	Description
Message	Displays the message type.
Received	Indicates the number of valid MC-Ring signaling messages received from the peer.
Transmitted	Indicates the number of valid MC-Ring signaling messages transmitted from the peer.
MCS ID Request	Displays the number of valid MCS ID requests were received from the peer.
MCS ID Response	Displays the number of valid MCS ID responses were received from the peer.
Ring Exists Request	Displays the number of valid 'ring exists' requests were received from the peer.
Ring Exists Response	Displays the number of valid ring exists' responses were received from the peer.
Keepalive	Displays the number of valid MC-Ring control packets of type 'keepalive' were received from the peer.

Sample Output: show redundancy multi-chassis mc-ring global-statistics

```

*A:ALA-48>show>redundancy>multi-chassis# mc-ring global-statistics
=====
Global MC Ring statistics
=====

```

```

Rx                                     : 0
Rx Too Short                         : 0
Rx Wrong Authentication              : 0
Rx Invalid TLV                      : 0
Rx Incomplete                       : 0
Rx Unknown Type                     : 0
Rx Unknown Peer                     : 0
Rx Unknown Ring                     : 0
Rx Unknown Ring Node                : 0
Tx                                   : 36763
Tx No Buffer                         : 0
Tx Transmission Failed               : 0
Tx Unknown Destination               : 0
Missed Configuration Events          : 0
Missed BFD Events                    : 0
=====
*A:ALA-48>show>redundancy>multi-chassis#

```

Output Fields: show redundancy multi-chassis mc-ring ring-node

[Table 94](#) describes the output fields for the mc-ring ring-node command.

Table 94 **Output Fields: Show redundancy multi-chassis mc-ring ring-node**

Label	Description
Oper State	Displays the state of the connection verification (both local and remote). notProvisioned — Connection verification is not provisioned. configErr — Connection verification is provisioned but a configuration error prevents it from operating properly. notTested — Connection verification is administratively disabled or is not possible in the current situation. testing — Connection Verification is active, but no results are yet available. connected — The ring node is reachable. disconnected — Connection verification has timed out.
In Use	Displays “True” if the ring node is referenced on an e-pipe or as an inter-dest-id on a static host or dynamic lease.

Output Fields: show redundancy multi-chassis mc-ring global-statistics

[Table 95](#) describes the output fields for the **show redundancy multi-chassis mc-ring global statistics** command.

Table 95 Output Fields: show redundancy multi-chassis mc-ring global-statistics

Label	Description
Rx	Displays the number of MC-ring signaling packets were received by this system.
Rx Too Short	Displays the number of MC-ring signaling packets were received by this system that were too short.
Rx Wrong Authentication	Displays the number of MC-ring signaling packets were received by this system with invalid authentication.
Rx Invalid TLV	Displays the number of MC-ring signaling packets were received by this system with invalid TLV.
Rx Incomplete	Displays the number of MC-ring signaling packets were received by this system that were incomplete.
Rx Unknown Type	Displays the number of MC-ring signaling packets were received by this system that were of unknown type.
Rx Unknown Peer	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown peer.
Rx Unknown Ring	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown ring.
Rx Unknown Ring Node	Displays the number of MC-ring signaling packets were received by this system that were related to an unknown ring node.
Tx	Displays the number of MC-ring signaling packets were transmitted by this system.
Tx No Buffer	Displays the number of MC-ring signaling packets could not be transmitted by this system due to a lack of packet buffers.
Tx Transmission Failed	Displays the number of MC-ring signaling packets could not be transmitted by this system due to a transmission failure.
Tx Unknown Destination	Displays the number of MC-ring 'unknown destination' signaling packets were transmitted by this system.
Missed Configuration Events	Displays the number of missed configuration events on this system.
Missed BFD Events	Displays the number of missed BFD events on this system.

sync

Syntax	sync [peer <i>ip-address</i>] sync peer <i>ip-address</i> detail sync [peer <i>ip-address</i>] statistics
Context	show>redundancy>multi-chassis
Description	This command displays synchronization information.
Parameters	<i>ip-address</i> — Specifies the IP address.
Values	ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF] H • d: [0 to 255] D
	detail — Displays detailed information.
	statistics — Displays statistics.

wpp

Syntax	wpp peer [<i>ip-address</i>] statistics
Context	show>redundancy>multi-chassis
Description	This command displays multi-chassis WPP information.
Parameters	<i>ip-address</i> — Specifies the IP address.
Values	ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF] H • d: [0 to 255] D
	statistics — Displays statistics.

ethernet

Syntax	ethernet
Context	show>system

Description This command enables the context to display system-specific Ethernet information.

efm-oam

Syntax **efm-oam**

Context show>system>ethernet

Description This command displays a system-level summary for EFM-OAM enabled ports. It includes the configuration, action, various states, and soft reset precedence for the protocol.

Output The following output is an example of EFM-OAM summary information.

Sample Output

```
show system ethernet efm-oam
=====
System Configuration
=====
Grace Tx Enable           : False
Dying Gasp On Reset       : Disabled
=====
Active Port Configuration & State Summary
=====
Fault Header:
  LinkMonitor: F = Frame, FP = Period, FS = Summary, SP = Symbol
  LocalSfAct  : CE = Critical, DG = Dying Gasp, EV = Event, PT = Port
  PeerRDIRx   : CE = Critical, DG = Dying Gasp, LF = Link, EV = Event
  Table Legend
    - = Not Configured, c = Configured, * = SF Active, s = Shutdown, l = Log Only
Soft Reset Header:
  Gasp : T = True, F = False (A = active, I = inactive)
  Grace: T = True, F = False (A = active, I = inactive)
  SR (Action on Soft Reset): - = None, DG = Dying Gasp, G = Grace TLV
TF (Trigger-Fault):
  - = Not configured, DG = Dying Gasp, CE = Critical Event
=====
Port          States  LinkMonitor LocalSfAct  PeerRdiRx  SoftReset  TF
              EFM LM  F FP FS SP  CE DG EV PT  CE DG LF EV  Gasp Grace SR
-----
1/1/1         Up  Dn  - - - -  - - - s  s s s l  T(I) T(I) - -
1/1/6         Dn  Dn  - - - -  - - - s  s s s l  T(I) T(I) - -
-----
No. of ports EFM enabled: 2
=====
```

license

Syntax **license available-licenses**

Context show>system

Description Displays information related to the active license(s) within the system. This includes the license name, target UUID, target product, target SR OS release, and the start, end, and issue date of the license.

If the **available-licenses** attribute is included, then the information for each active license shall be displayed. The usual case for more than one license in a system is during an upgrade procedures when there will needs to be both the originating and destination release licenses available in the system.

Parameters **available-licenses** — Includes information for all licenses active in the system.

Output The following output is an example of license information.

Sample Output

```
*A:ExitLeaf-35# show system license
=====
Current License
=====
License status : monitoring, valid license record
Time remaining : 131 days 9 hours
-----
License name   : sr-regress@list.nokia.com
License uuid   : ab516e50-2413-44aa-9f7c-34b4e5b64d19
Machine uuid   : ab516e50-2413-44aa-9f7c-34b4e5b64d19
License desc   : 7xxx Platform
License prod   : 7xxx Platform
License sros   : TiMOS-[BC]-16.0.*
Current date   : FRI NOV 03 15:53:54 UTC 2017
Issue date    : FRI SEP 22 20:55:14 UTC 2017
Start date    : FRI SEP 15 00:00:00 UTC 2017
End date      : THU MAR 15 00:00:00 UTC 2018
-----
Entitlement
IOM Upgrades  : !cr-er                                1
               : !er-he                                1
               : !cr400g-cr1200g+                       1
               : !er400g-er1200g+                       1
               : !he400g-he1200g+                       1
               : !cr-er1200g+                           1
               : !er-he1200g+                           1
MDA Upgrades  : cr1200g-cr1600g                         1
               : cr1600g-cr2400g                         1
               : er1200g-er1600g                         1
               : er1600g-er2400g                         1
               : he1200g-he1600g                         1
               : he1600g-he2400g                         1
               : cr1200g-er1200g                         1
               : er1200g-he1200g                         1
               : cr1600g-er1600g                         1
               : er1600g-he1600g                         1
               : cr2400g-er2400g                         1
               : er2400g-he2400g                         1
               : !any2400g-2400g+                       1
               : !cr3600g-cr4800g                       1
               : !er3600g-er4800g                       1
               : !he3600g-he4800g                       1
```

```

: !cr3600g-er3600g 1
: !er3600g-he3600g 1
: !cr4800g-er4800g 1
: !er4800g-he4800g 1
: !any3600g-3600gdd 1
: !any4800g-4800g+ 1
-----
! - indicates license not applicable to system

```

lldp

Syntax	lldp lldp neighbor
Context	show>system
Description	This command displays local Link Layer Discovery Protocol (LLDP) information at the system level. This includes an option to display summary information for all known peers.
Parameters	neighbor — Displays all LLDP neighbor information.
Output	The following output is an example of local LLDP information.

Sample Output: show system lldp

```

show system lldp
=====
LLDP Configuration
=====
Transmit Interval      : 30
Hold Multiplier        : 4
Reinit Delay           : 2
Notification Interval  : 5
Tx Credit Max          : 5
Message Fast Tx        : 1
Message Fast Tx Init   : 4
Admin Enabled          : True

-----
LLDP System Information
-----
Chassis Id Subtype     : 4
Chassis Id             : d8:1f:ff:00:00:00
System Name            : cses-V31
System Description     : TiMOS-B-0.0.I4269 both/i386 Nokia 7750 SR Copyright
                        : (c) 2000-2016 Nokia.
                        : All rights reserved. All use subject to applicable
                        : license agreements.
                        : Built on Wed Dec 3 19:14:27 PST 2014 by builder in /
                        : rel0.0/I4269/panos/main
Capabilities Supported  : bridge router
Capabilities Enabled    : bridge router
-----

```

```

LLDP Destination Addresses
-----
Index 1          : 01:80:c2:00:00:0e
Index 2          : 01:80:c2:00:00:03
Index 3          : 01:80:c2:00:00:00
-----

LLDP Remote Statistics
-----
Last Change Time   : 12/08/2014 21:34:48
Rem Table Inserts  : 10
Rem Table Deletes  : 1
Rem Table Drops    : 0
Rem Table Ageouts  : 3
-----

LLDP System Management Addresses
-----
Address SubType    : 1 (IPv4)
Address            : 10.1.1.31
Address If SubType : 2
Address If Id      : 1
Address OID        : .1.3.6.1.4.1.6527.1.3.3
Address SubType    : 2 (IPv6)
Address            : 2001:db8:dead:beef::31
Address If SubType : 2
Address If Id      : 1
Address OID        : .1.3.6.1.4.1.6527.1.3.3
-----

=====

show system lldp neighbor

Link Layer Discovery Protocol (LLDP) System Information
=====
NB = nearest-bridge   NTPMR = nearest-non-tpmr   NC = nearest-customer
=====
Lcl Port  Scope  Remote Chassis ID  Index  Remote Port  Remote System Name
-----
1/1/4     NB     D8:1D:FF:00:00:00  1      1/1/5        cses-v29
1/1/6     NB     D8:1D:FF:00:00:00  2      1/1/7        cses-v29
1/2/2     NB     D8:1F:FF:00:00:00  3      1/1/1        cses-v31
1/2/1     NB     D8:1E:FF:00:00:00  4      1/1/1        cses-v30
1/2/3     NB     D8:20:FF:00:00:00  5      1/1/1        cses-v32
1/2/4     NB     D8:21:FF:00:00:00  6      1/1/1        cses-V33
=====

Number of neighbors : 6

```

switch-fabric

Syntax **switch-fabric [exclude-sfm *sfm-list*]**
switch-fabric high-bandwidth-multicast

Context show>system

- Description

This command displays switch fabric information. When an MDA is not provisioned in an XCM, no information is displayed relating to that MDA. In order to display the forwarding capacity and multicast planes corresponding to that MDA, regardless of whether an XMA is equipped, an **mda-type** must be provisioned for that MDA.
- Parameters

sfm-list

— Specifies which SFMs to exclude from the displayed information.

Values

“A”, “B”, “A,B”, “1”, “2”, and “1,2”

high-bandwidth-multicast

— Displays MDA information about the switch-fabric plane’s high bandwidth multicast traffic tap allocation.
- Output

The following output is an example of switch fabric information, and [Table 96](#) describes switch-fabric output fields for 12-slot and 7-slot chassis models.

Sample Output

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:          Planes:
Slot/FP  Min  Max  Hbm Grp  Hi | Lo
-----
2/1      100% 100% No  0    1 0 2 3 4 5 6 7 8 9 10 11 12 13 14 | 15
5/1      100% 100% No  0    17 16 18 19 20 21 22 23 24 25 26 27 28 29 30 | 31
A         100% 100% No  0    40 | 40
B         100% 100% No  0    24 | 24
=====
*A:PE-1#
```

Table 96 Switch fabric output

Label	Description
Slot/FP	Displays the fabric slot within a chassis in the system. The CPM cards and XCM/IOM cards cannot be physically inserted into the switch fabric card slots.
Cap. Min.	Displays the minimum forwarding capacity of the slot and XMA/MDA as a percentage.
Cap. Max.	Displays the maximum forwarding capacity of the slot and XMA/MDA as a percentage.

2.21.2.5 Multilink Bundle Show Commands

The following multilink bundle show commands apply to the 7450 ESS and 7750 SR only:

multilink-bundle

Syntax	multilink-bundle [<i>bundle-id</i> <i>bgrp-id</i> <i>slot/mda</i> type { mlppp ima-grp mlfr }] [detail] multilink-bundle { <i>bundle-id</i> <i>bgrp-id</i> <i>slot/mda</i> } [ppp ima mlfr] multilink-bundle { <i>bundle-id</i> <i>bgrp-id</i> } relations
Context	show
Description	This command displays multilink bundle information. An operator can display: • All bundles on the system/MDA or all bundles of a given type on the system by specifying the value of type filter to be either mlppp, mlfr or ima-grp. • Bundle specific information in summary (no detail option) or detailed format (detail option specified) for one or more bundles • Protocol specific information (example PPP or IMA) for the specified bundle • ATM interface information for IMA groups see show port atm command for more details
Parameters	<i>bundle-id</i> — Specifies the multilink (PPP, MLFR or IMA) bundle to be associated with this IP interface. The command syntax must be used as follows: Syntax: <i>bundle-type-slot/mda.bundle-num</i> bundle: keyword type: ima, fr, ppp bundle-ppp-slot/mda.bundle-num (Creates a multilink PPP bundle.) bundle-ima-slot/mda.bundle-num (Creates an IMA group bundle.) bundle-fr-slot/mda.bundle-num (Creates an MLFR group bundle.) bundle-num: 1 to 336 <i>bgrp-id</i> — Specifies a bundle protection group. Values <i>bgrp-type-bgrp-num</i> bgrp: keyword type: ima, ppp bgrp-num: 1 to 2000 <i>slot/mda</i> — Specifies the slot and mda numbers associated to the IP interface. Values <i>slot/mda</i> slot: 1 to 10 mda: 1,2 type — Specifies the multi-link bundle group type. ppp — Displays PPP bundle information. ima, ima-grp — Displays IMA-type groups. mlfr — Displays bundle MLFR information, or used to display MLFR-type groups. mlppp — Displays MLPPP-type groups. detail — Provides detailed information.

relations — Displays the working and protection bundles associated with this bundle-id.

Output See the following sections for sample outputs:

- [Sample Output: show multilink-bundle](#)
- [Sample Output: show multilink-bundle detail](#)
- [Sample Output: show multilink-bundle ima-grp](#)
- [Sample Output: show multilink-bundle <bundle-id> \(showing PPP\)](#)
- [Sample Output: show multilink-bundle <bundle-id> detail \(showing Frame Relay\)](#)
- [Output Fields: show multilink-bundle](#)

Sample Output: show multilink-bundle

```
A:timetra-sim110# show multilink-bundle
=====
Bundle Summary
=====
Bundle          Type      Admin   Oper    Port      Min   Total/
Id              State    State   State   State     Links Active Links
-----
bundle-ppp-1/1.1  mlppp   Down    Down    Ghost      1     0/0
bundle-ima-1/1.2   ima     Down    Down    Link Up    1     1/0
-----
Bundles : 2
=====
A:timetra-sim110#
```

Sample Output: show multilink-bundle detail

```
*A:timetra-44-cpm2# show multilink-bundle detail
=====
Description      : MultiLink Bundle
Bundle Id        : bundle-ppp-1/1.1   Type           : mlppp
Admin Status     : up                Oper Status     : down
Minimum Links    : 1                Bundle IfIndex  : 572522497
Total Links      : 3                Active Links    : 0
Red Diff Delay   : 0                Yellow Diff Delay : 0
Red Diff Delay Act : none            MRRU           : 1524
Short Sequence   : true             Oper MRRU       : 1524
Oper MTU         : 1526             Fragment Threshold : 128 bytes
Up Time          : N/A              Bandwidth       : 0 KBit
PPP Input Discards : 0              Primary Member Port: 1/1/1.1.1.1.1.1
Mode             : access
Interleave-Frag  : false
=====
Member Port Id      #TS   Admin  Oper  Act  Down Reason      Up Time
-----
1/1/1.1.1.1.1.1     12    up     up    no   oper down        N/A
1/1/1.1.1.1.1.2     12    up     up    no   oper down        N/A
1/1/1.1.1.1.1.3     12    up     up    no   oper down        N/A
=====
Traffic Statistics
=====
Input              Output
```



```

-----
Octets                                0                                0
Packets                              0                                0
Errors                               0                                0
=====
Port Statistics
=====
                                     Input                                Output
-----
Packets                              0                                0
Discards                             0                                0
Unknown Proto Discards               0
=====
*A:timetra-44-cpm2#

```

Sample Output: show multilink-bundle ima-grp

```

A:timetra-sim110# show multilink-bundle type ima-grp
=====
Bundle Summary
=====
Bundle      Type      Admin      Oper      Port      Min      Total/
Id          State      State      State      Links    Active Links
-----
bundle-ima-1/1.2  ima      Down      Down      Link Up   1        1/0
-----
Bundles : 1
=====
A:timetra-sim110#

```

Sample Output: show multilink-bundle <bundle-id> (showing PPP)

```

A:timetra-sim110# show multilink-bundle bundle-ppp-1/1.1
=====
Bundle Summary
=====
Bundle      Type      Admin      Oper      Port      Min      Total/
Id          State      State      State      State      Links    Active Links
-----
bundle-ppp-1/1.1  mlppp    Down      Down      Ghost      1        0/0
-----
Bundles : 1
=====
A:timetra-sim110#

```

```

A:timetra-sim110# show multilink-bundle bundle-ppp-1/1.1 detail
=====
Bundle bundle-ppp-1/1.1 Detail
=====
Description      : MultiLink Bundle
Bundle Id        : bundle-ppp-1/1.1  Type           : mlppp
Admin Status     : down           Oper Status    : down
Minimum Links    : 1             Bundle IfIndex : 555745281
Total Links      : 0             Active Links   : 0
Red Diff Delay   : 0             Yellow Diff Delay : 0

```

```

Red Diff Delay Act : none           MRRU           : 1524
Short Sequence     : false          Oper MRRU       : 1524
Oper MTU           : 1522           Fragment Threshold : 128 bytes
Up Time            : N/A            Bandwidth        : 0 KBit
PPP Input Discards : 0              Primary Member Port: None
Interleave-Frag    : false

=====
Traffic Statistics
=====
                                     Input           Output
-----
Octets              0                   0
Packets             0                   0
Errors              0                   0
=====
Port Statistics
=====
                                     Input           Output
-----
Unicast Packets     0                   0
Multicast Packets   0                   0
Broadcast Packets   0                   0
Discards            0                   0
Known Proto Discards 0
=====
A:timetra-sim110#

```

Sample Output: show multilink-bundle <bundle-id> detail (showing Frame Relay)

```

*A:Cpm-A>config>port# show multilink-bundle

=====
Bundle Summary
=====
Bundle      Type      Admin   Oper   Port   Min   Total/
Id          State      State   State  State  Links Active Links
-----
bundle-fr-1/1.1  mlfr    Down    Down    Ghost    1     0/0
-----
Bundles : 1
=====
Bundle Summary
=====
Bundle      Type      Admin   Oper   Port   Min   Total/
Id          State      State   State  State  Links Active Links
-----
bundle-fr-1/1.1  mlfr    Down    Down    Ghost    1     0/0
-----
Bundles : 1
=====
*A:Cpm-A> show multilink-bundle bundle-fr-1/1.1 detail
=====
Bundle bundle-fr-1/1.1 Detail
=====
Description      : MultiLink Bundle
Bundle Id        : bundle-fr-1/1.1      Type           : mlfr
Admin Status     : down              Oper Status     : down
Minimum Links    : 1              Bundle IfIndex  : 572530689

```

```

Total Links      : 0
Red Diff Delay   : 0
Red Diff Delay Act : none
Short Sequence   : N/A
Oper MTU         : 0
Up Time          : N/A
PPP Input Discards : N/A
Mode             : access
Interleave-Frag  : N/A

Active Links     : 0
Yellow Diff Delay : 0
MRRU            : N/A
Oper MRRU       : N/A
Fragment Threshold : 128 bytes
Bandwidth       : 0 KBit
Primary Member Port: None

```

=====
Traffic Statistics
=====

	Input	Output
Octets	0	0
Packets	0	0
Errors	0	0

=====
Port Statistics
=====

Input	Output
Unicast Packets	0
Multicast Packets	0
Broadcast Packets	0
Discards	0
Unknown Proto Discards	0

*A:Cpm-A> show multilink-bundle bundle-fr-1/1.1 mlfr frame-relay

=====
Frame Relay Info for bundle-fr-1/1.1
=====

```

Mode           : dte          LMI Type           : itu
FR Interface Status : fault
N391 DTE       : 6            N392 DCE           : 3
N392 DTE       : 3            N393 DCE           : 4
N393 DTE       : 4            T392 DCE           : 15
T391 DTE       : 10
Tx Status Enquiry : 0          Rx Status Enquiry  : 0
Rx Status Messages : 0          Tx Status Messages : 0
Status Message Timeouts : 0      Status Enquiry Timeouts : 0
Discarded Messages : 0          Inv. RxSeqNum Messages : 0

```

=====
Service Access Points(SAP)
=====

```

Service Id      : 39
SAP             : 1/1/2.3.5.2.2:18      Encap           : frRel
Description     : Default sap description for service id 39
Admin State     : Up                    Oper State      : Up
Flags           : None
Multi Svc Site  : None
Last Status Change : 12/02/2008 20:48:17
Last Mgmt Change  : 12/02/2008 20:46:38
Sub Type        : regular
Split Horizon Group: (Not Specified)

```

Admin MTU	: 4474	Oper MTU	: 4474
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id	: n/a
Ingr Mac Fltr-Id	: n/a	Egr Mac Fltr-Id	: n/a
Ingr IPv6 Fltr-Id	: n/a	Egr IPv6 Fltr-Id	: n/a
tod-suite	: None	qinq-pbit-marking	: both
Ing Agg Rate Limit	: max	Egr Agg Rate Limit	: max
Endpoint	: N/A		

FRF-12	: Disabled	Collect Stats	: Disabled
Acct. Pol	: None		

FRF12 on channel where sap resides:

=====

Service Access Points(SAP)

=====

Service Id	: 1		
SAP	: 1/1/1.3.7.4.1:16	Encap	: frRel
Description	: sap-1-88.10.131.1		
Admin State	: Up	Oper State	: Up
Flags	: None		
Multi Svc Site	: None		
Last Status Change	: 12/02/2008 20:48:15		
Last Mgmt Change	: 12/02/2008 20:46:36		
Sub Type	: regular		
Split Horizon Group	: (Not Specified)		
Admin MTU	: 9194	Oper MTU	: 9194
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id	: n/a
Ingr Mac Fltr-Id	: n/a	Egr Mac Fltr-Id	: n/a
Ingr IPv6 Fltr-Id	: n/a	Egr IPv6 Fltr-Id	: n/a
tod-suite	: None	qinq-pbit-marking	: both
Ing Agg Rate Limit	: max	Egr Agg Rate Limit	: max

FRF-12 (I/F)	: Enabled		
Scheduling Class	: 0		
Acct. Pol	: None	Collect Stats	: Disabled
Anti Spoofing	: None	Avl Static Hosts	: 0
		Tot Static Hosts	: 0

Calling-Station-Id : n/a
Application Profile: None

FRF12 ETE on sap capable of supporting it:

=====

Service Access Points(SAP)

=====

Service Id	: 1		
SAP	: 1/1/1.1:16	Encap	: frRel
Description	: (Not Specified)		
Admin State	: Up	Oper State	: Up
Flags	: None		
Multi Svc Site	: None		
Last Status Change	: 12/02/2008 20:48:12		
Last Mgmt Change	: 12/02/2008 20:46:36		
Sub Type	: regular		
Split Horizon Group	: (Not Specified)		

```

Admin MTU           : 9194
Ingr IP Fltr-Id     : n/a
Ingr Mac Fltr-Id    : n/a
Ingr IPv6 Fltr-Id   : n/a
tod-suite           : None
Ing Agg Rate Limit  : max

FRF-12 (ETE)        : Enabled
Scheduling Class    : 3
Acct. Pol           : None
Disabled
Anti Spoofing       : None

Calling-Station-Id  : n/a
Application Profile  : None
-----
Oper MTU            : 9194
Egr IP Fltr-Id     : n/a
Egr Mac Fltr-Id    : n/a
Egr IPv6 Fltr-Id   : n/a
qinq-pbit-marking   : both
Egr Agg Rate Limit : max

Ete-Frag-Threshold : 128
Collect Stats       :
Avl Static Hosts    : 0
Tot Static Hosts    : 0

```

Output Fields: show multilink-bundle

[Table 97](#) describes the output fields for the **show multilink-bundle** command.

Table 97 Output Fields: show multilink-bundle

Label	Description
Type	Specifies the type of this multilink bundle. mlppp — Indicates that the bundle is of type MLPPP. ima — Indicates that the bundle is of type IMA group.
Admin State	Up — The bundle is administratively up. Down — The bundle is administratively down.
Oper State	Up — The bundle is operationally up. Down — The bundle is operationally down.
Port State	Displays the state level of the port. none — Indicates that the port is either in its initial creation state or is just about to be deleted. ghost — No member links are configured as part of this bundle. down — All member links are in “none”, “ghost”, or “down” state. linkUp — At least one member link is in port state “link up” but the bundle protocol is not yet operationally up (due to bundle protocol still coming up. For example, due to insufficient number of member links in “link up” state yet or to bundle being shut down. Up — Indicates that the bundle is ready to pass some kinds of traffic as the bundle protocol has come up (at least “minimum links” member links are in the port state up and the bundle protocol is up.)
Min Links Minimum Links	Displays the minimum number of links that must be active for a bundle to be active. If the number of active links drop below the given minimum then the multilink bundle will transition to an operation down state.

Table 97 **Output Fields: show multilink-bundle (Continued)**

Label	Description (Continued)
Bundle IfIndex	Displays the bundle's interface index number which reflects its initialization sequence.
Total Links	Displays the total number of member links configured for this bundle.
Active Links	Displays the total number of active links for the bundle.
Description	Displays configured description for this bundle.
Bundle Id	Displays the port ID for this bundle.
Red Diff Delay	Displays the maximum acceptable differential delay for individual circuits within this multilink bundle. If the delay exceeds this threshold, a trap is issued. The differential delay is calculated as the round-trip differential delay for MLPPP bundles, and as uni-directional differential delay for IMA bundles.
Fragment Threshold	Displays configured fragment threshold value for this bundle.
Up Time	Displays time elapsed since the last bundle transition to Up when part of bundle information. Displays time elapsed since the last link transition to active when part of member information.
Bandwidth	Displays bandwidth configured for this IMA bundle in kbytes.
Primary Port Member	Displays the portId of the IMA group member chosen as a Primary Port.
Member Port Id	Displays portId of each member of this bundle.
Admin	Displays administrative port status of a member link.
Oper	Displays operational port status of a member link.
Active	Displays whether a member link is active or not.
Down Reason	Displays the reason for why a member link is not active.
Traffic and Port statistics	The traffic and port statistics information displayed for bundles when detail option is selected is the same as information displayed for physical ports.

relations

Syntax **relations**

Context show>multilink-bundle

Description This command displays the working and protection bundles associated with this bundle-id.

Output The following output is an example of multilink bundle relations, and [Table 98](#) describes the output fields.

Sample Output: show multilink-bundle <bundle-id> relations

```
A:ALA-48>show# show multilink-bundle bundle-ima-1/1.2 relations
=====
Bundle Relationship
=====
Bundle      Admin  Oper   Working      Protect      Active
Id          State  State  Bundle Id    Bundle Id     Bundle
-----
bpgrp-ima-1  Down   Down   bundle-ima-1/1.1  bundle-ima-1/1.2  Protect
-----
Bundles : 1
=====
A:ALA-48>show#
```

Table 98 Output Fields: show multilink-bundle relations

Label	Description
BundleID	Displays the bundle number.
Admin State	Up—The bundle is administratively up. Down — The bundle is administratively down.
Oper State	Up — The bundle is operationally up. Down—The bundle is operationally down.
Working BundleID	Displays the bundle that is currently in working mode.
Protect BundleID	Displays the bundle that is currently in protect mode.
Active Bundle	Displays the mode of the active bundle.

ima

Syntax	ima
Context	show>multilink-bundle
Description	This command enables the context to display IMA group data.
Output	The following output is an example of IMA information, and Table 99 describes the output fields.

Sample Output: show multilink-bundle <bundle-id> ima

```
A:timetra-sim110# show multilink-bundle bundle-ima-1/1.2
=====
Bundle Summary
=====
Bundle      Type   Admin   Oper   Port   Min   Total/
Id          State  State  State  State  Links Active Links
-----
```

```

-----
bundle-ima-1/1.2      ima      Down      Down      Link Up    1      1/0
-----
Bundles : 1
=====
A:timetra-sim110#

A:timetra-sim110# show multilink-bundle bundle-ima-1/1.2 detail
=====
Bundle bundle-ima-1/1.2 Detail
=====
Description          : MultiLink Bundle
Bundle Id             : bundle-ima-1/1.2   Type              : ima
Admin Status         : down               Oper Status       : down
Minimum Links        : 1                  Bundle IfIndex    : 555749378
Total Links          : 1                  Active Links      : 0
Red Diff Delay       : 25                  Yellow Diff Delay : N/A
Red Diff Delay Act   : down               MRRU              : N/A
Short Sequence       : N/A                Oper MRRU         : N/A
Oper MTU             : 1524               Fragment Threshold : 128 bytes
Up Time              : N/A                Bandwidth         : 0 KBit
PPP Input Discards   : N/A                Primary Member Port: 1/1/1.1.1.1
Interleave-Frag      : N/A
-----
Member Port Id      Admin Oper Active Down Reason      Up Time
-----
1/1/1.1.1.1        up    up    no    oper down        N/A
=====
Traffic Statistics
=====
                                     Input      Output
-----
Octets              0              0
Packets             0              0
Errors              0              0
=====
Port Statistics
=====
                                     Input      Output
-----
Packets             0              0
Discards            0              0
Unknown Proto Discards 0
=====
A:timetra-sim110#

A:timetra-sim110# show multilink-bundle bundle-ima-1/1.2 ima
=====
Bundle bundle-ima-1/1.2 IMA group information
=====
Version             : 1.1
Current State       : Startup Near-end
Near-end State      : Startup
Far-end State       : Not configured
Group Test State    : Disabled
Max BW Links        : 8
Operational Secs    : N/A                Down Secs         : 2811
Tx IMA Id           : 0                  Rx IMA Id         : 255
Tx Timing Ref Link  : N/A                Rx Timing Ref Link : N/A

```



```

Tx Oam Label      : 3                Rx Oam Label      : 0
Test Link         : N/A              Test Pattern     : 0
Near-End Clock-Mode: ctc              Far-End Clock-Mode: itc
Link Deact Timer  : 2000              Link Act Timer   : 10000
Alpha-value       : 2                Beta-value      : 2
Gamma-value       : 1                Symmetry Mode   : symmetric
Tx CR Available   : 0 KBit            Rx CR Available  : 0 KBit
Least Delayed Link: N/A              Max Obs Diff Delay: 0
Near-End Fails    : 1                Far-end Fails    : 0
Tx Icp Cells      : 0                Rx Icp Cells     : 0
Errored Icp Cells : 0                Rx Lost Icp Cells: 0
=====
A:timetra-sim110#

```

The following stats should only be displayed if the bundle type is mlppp-lfi and the detail is issued.

```

LFI Statistics
IngressEgress
Packet Bytes PacketByte
High500050000000200020000000
Normal100001000000050005000000
=====

```

Table 99 **Output Fields: show multilink-bundle bundle-ima**

Label	Description
BundleId	Displays the bundle ID number.
Type	Specifies the type of this multilink bundle. mlppp — Indicates that the bundle is of type MLPPP. ima — Indicates that the bundle is of type IMA group.
Admin State	ima — Indicates that the bundle is of type IMA group. Down — The bundle is administratively down.
Oper State	Up — The bundle is operationally up. Down — The bundle is operationally down.

Table 99 Output Fields: show multilink-bundle bundle-ima (Continued)

Label	Description (Continued)
Port State	Displays the state level of the port. None — Indicates that the port is either in its initial creation state or is just about to be deleted. Ghost — No member links are configured as part of this bundle. Down — All member links are in “none”, “ghost”, or “down” state. linkUp — At least one member link is in port state “link up” but the bundle protocol is not yet operationally up (due to bundle protocol still coming up. For example, due to insufficient number of member links in “link up” state yet or to bundle being shut down. Up — Indicates that the bundle is ready to pass some kinds of traffic as the bundle protocol has come up (at least “minimum links” member links are in the port state up and the bundle protocol is up).
Min Links Minimum Links	Displays the minimum number of links that must be active for a bundle to be active. If the number of links drop below the given minimum then the multilink bundle will transition to an operation down state).
Total/Active Links	Displays the total number of active links for the bundle.

ppp

Syntax	ppp [multiclass]
Context	show>multilink-bundle
Description	This command enables the context to display PPP group data.
Parameters	multiclass — Specifies to display multi-class MLPPP information.
Output	The following output is an example of PPP information.

Sample Output: show multilink-bundle <bundle-id> ppp

```
A:timetra-sim110# show multilink-bundle bundle-ppp-1/1.1 ppp
=====
PPP Protocols for bundle-ppp-1/1.1
=====
Protocol  State      Last Change      Restart Count    Last Cleared
-----
ipcp      initial    02/16/2007 06:11:44      0      02/16/2007 06:11:44
mplscp    initial    02/16/2007 06:11:44      0      02/16/2007 06:11:44
bcp       initial    02/16/2007 06:11:44      0      02/16/2007 06:11:44
osicp     initial    02/16/2007 06:11:44      0      02/16/2007 06:11:44
ipvc6cp   initial    02/16/2007 06:11:44      0      02/16/2007 06:11:44
```

```
=====
Local Mac address   : 8c:6e:01:01:00:3d  Remote Mac address : 00:00:00:00:00:00
Local IPv4 address  : 0.0.0.0             Remote IPv4 address: 0.0.0.0
Local IPv6 address  : ::
Remote IPv6 address: ::
=====
*A:mlppp_top#

*A:mlppp_top# show multilink-bundle bundle-ppp-1/1.1 ppp multiclass
=====
MLPPP Per Class Traffic Statistics for bundle-ppp-1/1.1
=====
                                     Input          Output
-----
Class 0
  Octets                0                      0
  Packets               0                      0
  Errors                0                      0
Class 1
  Octets                0                      0
  Packets               0                      0
  Errors                0                      0
Class 2
  Packets               0                      0
  Errors                0                      0
Class 3
  Octets                0                    270400
  Packets               0                    2704
  Errors                0                      0
=====
*A:mlppp_top#
```

atm

Syntax	atm [detail]
Context	show>ml-bundle>ima
Description	This command displays multilink bundle IMA ATM information.
Parameters	detail — Displays detailed information.
Output	The following output is an example of IMA ATM information, and Table 100 describes the output fields.

Sample Output: show multilink-bundle <bundle-id> ima atm>

```
A:NS052651098# show multilink-bundle bundle-ima-1/1.1 ima atm
=====
ATM Info for bundle-ima-1/1.1
=====
Cell Mode           : UNI           Mapping           : Direct
Configured VCs      : 1             Configured VPs     : 0
Configured VTs      : 0             Configured IFCs    : 0
```

```

Configured minimum VPI: 0
Last Unknown VPI/VCI : none
=====
ATM Bandwidth Info
=====
              kbps      %              kbps      %
-----
Ingress CBR      : 15232    100%    Egress CBR      : 15232    100%
Ingress RT-VBR   : 0        0%      Egress RT-VBR   : 0        0%
Ingress NRT-VBR  : 0        0%      Egress NRT-VBR  : 0        0%
Ingress UBR      : 0        0%      Egress UBR      : 0        0%
-----
Ingress Total    : 15232    100%    Egress Total    : 15232    100%
ATM Link Bandwidth : 15232 kbps
Shaped Bandwidth  : 15232 kbps

```

Table 100 **Output Fields: show multilink-bundle <bundle-id> ima atm**

Label	Description
Cell Mode	Displays the cell format (UNI or NNI) used on the ATM interface.
Configured VCs	Displays the number of configured VCs.
Configured VTs	Displays the number of configured VTs.
Configured minimum VPI	Displays the minimum VPI configured for this bundle.
Last Unknown VPI/VCI	Indicates the last unknown VPI/VCI that was received on this interface.
Mapping	Displays ATM cell mapping used on this interface: Direct or PLCP.
Configured VPs	Displays the number of configured VPs.
Configured IFCs	Displays the number of configured IFCs.

connections

- Syntax** **connections**
- Context** show>ml-bundle>ima>atm
- Description** This command displays connection information.
- Output** The following output is an example of IMA ATM connections information, and [Table 101](#) describes the output fields.

Sample Output: show multilink-bundle bundle-ima ima atm connections

```
A:NS052651098# show multilink-bundle bundle-ima-1/1.1 ima atm connections
=====
ATM Connections, Port bundle-ima-1/1.1
=====
      Owner  Type      Ing.TD  Egr.TD  Adm  OAM      Opr
-----
1/100    SAP    PVC       2       2    up   up       up
=====
A:NS052651098#
```

Table 101 **Output Fields: show multilink-bundle <bundle-id> ima atm connections**

Label	Description
Owner	Identifies the system entity that owns a specific ATM connection.
Type	Specifies the type of this multilink bundle. mlppp — Indicates that the bundle is of type MLPPP. ima — Indicates that the bundle is of type IMA group.
Ing.TD	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr.TD	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Adm	ima — Indicates that the bundle is of type IMA group. Down — The bundle is administratively down.
Opr	Up — The bundle is operationally up. Down — The bundle is operationally down.
OAM	Indicates the OAM operational status of ATM connections.

port-connection

Syntax	port-connection [detail]
Context	show>ml-bundle>ima>atm
Description	This command displays port connection information.
Parameters	detail — Displays detailed information.
Output	The following output is an example of port connection information, and Table 102 describes the output fields.

Sample Output: show multilink-bundle <bundle-id> ima atm port-connection

```

A:NS052651098# show multilink-bundle bundle-ima-1/1.1 ima atm port-connection
=====
ATM Port Connection
=====
Port Id           : bundle-ima-1/1.1
Admin State       : up                Oper state        : up
Owner             : SAP
Endpoint Type     : Port              Cast Type         : P2P
Ing. Td Idx       : 2                 Egr. Td Idx       : 2
Last Changed      : 01/16/2007 14:24:00
=====
A:NS052651098#

```

Table 102 **Output Fields: show multilink-bundle <bundle-id> ima atm port-connection**

Label	Description
Port ID	Displays the port ID for this bundle.
Admin State	ima — Indicates that the bundle is of type IMA group. Down — The bundle is administratively down.
Oper State	Up — The bundle is operationally up. Down—The bundle is operationally down.
Owner	Identifies the system entity that owns a specific ATM connection.
Endpoint Type	Displays the endpoint type.
Cast Type	Indicates the connection topology type.
Ing. Td Idx	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr. Td Idx	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Last Changed	Indicates the date and time when the interface connection entered its current operational state.

pvc

Syntax **pvc** [*vpi/vci*] [**detail**]

Context show>ml-bundle>ima>atm

Description This command displays ATM port PVC information.

Parameters *vpi* — Displays the VPI values.

Values vpi:0 to 4095 (NNI)
 0 to 255 (UNI)

vci: — Displays the VCI values.

Values 1, 2, 5 to 65535

detail — Provides detailed information.

Output The following output is an example of PVC information, and [Table 103](#) describes the output fields.

Sample Output: show multilink-bundle <bundle-id> ima atm pvc

```
A:NS052651098# show multilink-bundle bundle-ima-1/1.1 ima atm pvc
=====
ATM PVCs, Port bundle-ima-1/1.1
=====
VPI/VCI   Owner  Type      Ing.TD  Egr.TD  Adm  OAM      Opr
-----
1/100     SAP    PVC        2        2        up   up       up
=====
A:NS052651098#
```

```
A:NS052651098# show multilink-bundle bundle-ima-1/1.1 ima atm pvc detail
=====
```

```
ATM PVCs, Port bundle-ima-1/1.1
=====
```

```
VPI/VCI   Owner  Type      Ing.TD  Egr.TD  Adm  OAM      Opr
-----
1/100     SAP    PVC        2        2        up   up       up
=====
```

```
=====
ATM Statistics
=====
```

	Input	Output
Octets	0	0
Cells	0	0

```
=====
AAL-5 Packet Statistics
=====
```

	Input	Output
Packets	0	0
Dropped Packets	0	0
CRC-32 Errors	0	
Reassembly Timeouts	0	
Over Sized SDUs	0	

```
=====
ATM OAM Statistics
=====
```

	Input	Output
AIS	0	0

```

RDI                                0                                0
Loopback                          0                                0
CRC-10 Errors                      0
Other                              0
=====
A:NS052651098#

```

Table 103 Output Fields: show multilink-bundle <bundle-id> ima atm pvc

Label	Description
VPI/VCI	Displays the VPI/VCI value.
Owner	Specifies the type of this multilink bundle. mlppp — Indicates that the bundle is of type MLPPP. ima — Indicates that the bundle is of type IMA group.
Type	Identifies the system entity that owns a specific ATM connection.
Ing.TD	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr.TD	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Adm	ima — Indicates that the bundle is of type IMA group. Down — The bundle is administratively down.
Opr	Up — The bundle is operationally up. Down — The bundle is operationally down.
OAM	Indicates the OAM operational status of ATM connections.

pvp

Syntax pvp [vpi] [detail]**Context** show>ml-bundle>ima>atm**Description** This command displays ATM port PVP information.**Parameters** vpi — Displays the VPI values.

Values vpi: 0 to 4095 (NNI)
0 to 255 (UNI)

detail — Displays detailed information.**Output** The following output is an example of PVP information, and [Table 104](#) describes the output fields.

Sample Output: show multilink-bundle <bundle-id> ima atm pvp

```
A:ima2# show multilink-bundle bundle-ima-1/1.1 ima atm pvp
=====
ATM PVPs, Port bundle-ima-1/1.1
=====
VPI      Owner  Type    Ing.TD  Egr.TD  Adm  OAM      Opr
-----
2        SAP    PVP     1       1       up   up       up
=====
A:ima2#
```

Table 104 **Output Fields: show multilink-bundle <bundle-id> ima atm pvp**

Label	Description
VPI	Displays the VPI value.
Owner	Identifies the system entity that owns a specific ATM connection.
Type	Specifies the type of this multilink bundle. mlppp — Indicates that the bundle is of type MLPPP. ima — Indicates that the bundle is of type IMA group.
Ing.TD	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Eng.TD	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Adm	ima — Indicates that the bundle is of type IMA group. Down — The bundle is administratively down.
OAM	Indicates the OAM operational status of ATM connections.
Opr	Up — The bundle is operationally up. Down — The bundle is operationally down.

pvt

Syntax **pvt** [*vpi.vpi*] [**detail**]

Context show>ml-bundle>ima>atm

Description This command displays ATM port PVT information.

Parameters *vpi.vpi* — Displays the VPI values.

Values vpi:0 to 4095 (NNI)
 0 to 255 (UNI)

vpi:0 to 4095 (NNI)
0 to 255 (UNI)

detail — Provides detailed information.

Output The following output is an example of PVT information, and [Table 105](#) describes the output fields.

Sample Output: show multilink-bundle <bundle-id> ima atm pvt

```
A:ima2# show multilink-bundle bundle-ima-1/1.1 ima atm pvt
=====
ATM PVTs, Port bundle-ima-1/1.1
=====
VPI Range  Owner  Type      Ing.TD  Egr.TD  Adm          Opr
-----
4.5         SAP    PVT       1       1       up           up
=====
A:ima2#
```

Table 105 **Output Fields: show multilink-bundle <bundle-id> ima atm pvt**

Label	Description
VPI Range	Displays the VPI range:
Owner	Identifies the system entity that owns a specific ATM connection.
Type	Specifies the type of this multilink bundle. mlppp — Indicates that the bundle is of type MLPPP. ima — Indicates that the bundle is of type IMA group.
Ing.Td	Specifies the ATM traffic descriptor profile that applies to the receive direction of the interface connection.
Egr.Td	Specifies the ATM traffic descriptor profile that applies to the transmit direction of the interface connection.
Adm	ima — Indicates that the bundle is of type IMA group. Down — The bundle is administratively down.
Opr	Up — The bundle is operationally up. Down — The bundle is operationally down.

2.21.2.6 LAG Show Commands

lag

Syntax	<pre> lag [<i>lag-id</i>] [detail] [statistics] lag [<i>lag-id</i>] description lag [<i>lag-id</i>] port lag <i>lag-id</i> associations per-link-hash interface [class {1 2 3}] lag <i>lag-id</i> associations lag <i>lag-id</i> bfd lag <i>lag-id</i> [detail] eth-cfm [tunnel <i>tunnel-id</i>] lag <i>lag-id</i> associations per-link-hash interface lag <i>lag-id</i> associations link-map-profile [<i>link-map-profile</i>] interface lag <i>lag-id</i> lACP-partner lag <i>lag-id</i> detail lACP-partner lag <i>lag-id</i> link-map-profile <i>link-map-profile</i> lag <i>lag-id</i> per-link-hash [class {1 2 3}] lag <i>lag-id</i> associations per-link-hash sap [class {1 2 3}] lag <i>lag-id</i> associations link-map-profile [<i>link-map-profile</i>] sap lag <i>lag-id</i> per-link-hash port <i>port-id</i> </pre>
Context	show
Description	This command displays Link Aggregation Group (LAG) information. If no command line options are specified, a summary listing of all LAGs is displayed.
Parameters	<i>lag-id</i> — Displays only information on the specified LAG ID. Values 1 to 800 Default Displays information for all LAG IDs. detail — Displays detailed LAG information. Default Displays summary information. statistics — Displays LAG statistics information. port — Displays port information. description — Displays the user-defined description for the LAG. associations — Displays a list of current router interfaces to which the LAG is assigned. interface — Displays interface information. bfd — Displays per-member BFD information. eth-cfm — Displays a list of Ethernet tunnels to which the LAG is assigned.

tunnel-id — Filters ETH-CFM MEP information to one tunnel.

Values 1 to 4094

per-link-hash — Displays information about a SAP or interface associated with this LAG will send traffic over a single link of a LAG auto-re-balancing as links are added and removed from this LAG.

lACP-partner — Displays LACP partner information.

link-map-profile — Displays information about a specified LAG link map profile.

Values 1 to 64

sap — Displays SAP information.

port-id — Specifies the port ID.

Values *slot/mda/port*

Output See the following sections for output samples:

- [Sample Output: show lag](#)
- [Output Fields: show lag](#)
- [Sample Output: show lag <lag-id> detail](#)
- [Output Fields: show lag <lag-id> detail](#)
- [Sample Output: Show LAG Statistics](#)
- [Output Fields: Show LAG Statistics](#)
- [Sample Output: Show LAG <lag-id> Associations](#)
- [Sample Output: Show LAG <lag-id> detail \(with MC-LAG Output\)](#)
- [Sample Output: show lag <lag-id> detail \(Without MC-LAG Output\)](#)
- [Sample Output: show lag <lag-id> lACP-partner](#)
- [Output Fields: Show LAG <lag-id> Associations](#)

Sample Output: show lag

```
A:ALA-48>config# show lag
=====
Lag Data
=====
```

Lag-id	Adm	Opr	Port-Threshold	Up-Link-Count	MC Act/Stdby
1	up	down	0	0	N/A
2	up	up	0	1	active
3	up	down	0	0	standby
4	up	down	0	0	standby
10	up	down	0	0	N/A

```
-----
Total Lag-ids: 5          Single Chassis: 2          MC Act: 1          MC Stdby: 2
=====
A:ALA-48>config# show lag
```

```
A:sr7- show lag 10 port
```

```
=====
Lag Port States
LACP Status: e - Enabled, d - Disabled
=====
Lag-id Port-id  Adm  Act/Stdby Opr   Primary  Sub-group   Forced  Priority
-----
10(e)  1/1/8    up   active   up    yes      1           -      32768
        1/1/9    up   standby  down           2           -      32768
=====
```

Output Fields: show lag

[Table 106](#) describes the output fields for the **show lag** command.

Table 106 Output Fields: show lag

Label	Description
LAG ID	The LAG or multi-link bundle ID that the port is assigned to.
Adm	Up — The LAG is administratively up. Down — The LAG is administratively down.
Opr	Up — The LAG is operationally up. Down—The LAG is operationally down.
Port-Threshold	The number of operational links for the LAG at or below which the configured action will be invoked.
Up-Link-Count	The number of ports that are physically present and have physical links present.
MC Act/Stdby	Member port is selected as active or standby link.

Sample Output: show lag <lag-id> detail

```
ALA-1# show lag 10 detail
=====
LAG Details
=====
Description          : N/A
-----
Details
-----
Lag-id           : 10                Mode           : network
Adm              : up                Opr            : up
Thres. Exceeded Cnt : 17              Port Threshold : 0
Thres. Last Cleared : 01/22/2000 19:41:38 Threshold Action : down
Dynamic Cost      : false            Encap Type     : null
Configured Address : 0c:a4:02:20:69:4b Lag-IfIndex    : 1342177290
Hardware Address   : 0c:a4:02:20:69:4b
Hold-time Down    : 0.0 sec          Port Type      : standard
Per FP Ing Queuing : disabled
LACP              : enabled          Mode           : active
```

```

LACP Transmit Intvl : fast           LACP xmit stdby      : enabled
Selection Criteria  : highest-count  Slave-to-partner     : disabled
MUX control         : coupled
Number of sub-groups: 2              Forced                : -
System Id           : 0c:a4:02:20:68:01 System Priority      : 32768
Admin Key           : 32770           Oper Key              : 32770
Prtr System Id      : 0c:a4:02:1f:88:01 Prtr System Priority  : 32768
Prtr Oper Key       : 32771
Standby Signaling   : lacp

```

```

-----
Port-id      Adm      Act/Stdby Opr      Primary  Sub-group  Forced  Prio
-----
1/1/8        up        active   up        yes       1          -      32768
1/1/9        up        standby  down               2          -      32768

```

```

-----
Port-id      Role      Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
-----
1/1/8        actor    No   No   Yes   Yes  Yes  Yes  Yes      Yes
1/1/8        partner No   No   Yes   Yes  Yes  Yes  Yes      Yes
1/1/9        actor    No   No   No    No   No   Yes  Yes      Yes
1/1/9        partner No   No   No    No   No   Yes  Yes      Yes

```

```
=====
ALA-1#
```

Output Fields: show lag <lag-id> detail

Table 107 describes the output fields for the **show lag lag-id detail** command. The output is dependent on whether or not the LAG was configured as a multi-chassis (MC) LAG.

Table 107 Output Fields: show lag <lag-id> detail

Label	Description
LAG ID	The LAG or multi-link trunk (MLT) that the port is assigned to.
Adm	Up — The LAG is administratively up. Down — The LAG is administratively down.
Port Threshold	If the number of available links is equal or below this number, the threshold action is executed.
Thres. Last Cleared	The last time that keepalive stats were cleared.
Dynamic Cost	The OSPF costing of a link aggregation group based on the available aggregated, operational bandwidth.
Configured Address	The configured Ethernet MAC address.
Hardware Address	The Ethernet MAC address.
Hold-Time Down	The timer, in tenths of seconds, which controls the delay between detecting that a LAG is down and reporting it to the higher levels.

Table 107 Output Fields: show lag <lag-id> detail (Continued)

Label	Description (Continued)
LACP	Enabled — LACP is enabled. Down — LACP is disabled.
LACP Transmit Intvl	LACP timeout signaled to peer.
Selection Criteria	Configured subgroup selection criteria.
MUX control	Configured type of multiplexing machine control used in a LAG with LACP in active/ passive modes. coupled — TX and RX activate together. independent — RX activates independent of TX.
Number of subgroups	Total subgroups in LAG.
System ID	System ID used by actor in LACP messages.
Admin Key	Configured LAG key.
Oper Key	Key used by actor in LACP messages.
System Priority	System priority used by actor in LACP messages.
Prtr System ID	System ID used by partner in LACP messages.
Prtr Oper Key	Key used by partner in LACP messages.
Prtr System Priority	System priority used by partner in LACP messages.
Mode	LAG in access or network mode.
Opr	Up — The LAG is operationally up. Down — The LAG is operationally down.
Port Threshold	Configured port threshold.
Thres. Exceeded Cnt	The number of times that the drop count was reached.
Threshold Action	Action to take when the number of available links is equal or below the port threshold.
Encap Type	The encapsulation method used to distinguish customer traffic on a LAG.
Lag-IFIndex	A box-wide unique number assigned to this interface.
Adapt QoS	Displays the configured QoS mode.
Port ID	The specific port ID.
(LACP) Mode	LACP active or passive mode.
LACP xmit standby	LACP transmits on standby links enabled / disabled.

Table 107 Output Fields: show lag <lag-id> detail (Continued)

Label	Description (Continued)
Slave-to-partner	Configured enabled/disabled.
Port-id	Displays the member port ID.
Adm	Displays the member port administrative state.
Active/stdby	Indicates that the member port is selected as the active or standby link.
Opr	Indicates that the member port operational state.
Primary	Indicates that the member port is the primary port of the LAG.
Sub-group	Displays the member subgroup where the member port belongs to.
Priority	Displays the member port priority.

Sample Output: Show LAG Statistics

```

=====
LAG Statistics
=====
-----
LAG 2
-----
Description          : N/A
Port-id      Input (Packets)      Output (Packets)
              Input (Bytes)       Output (Bytes)
              Input Errors        Output Errors
-----
1/1/4         0                  0
              0                  0
              0                  0
-----
Totals        0                  0
              0                  0
              0                  0
=====

```

Output Fields: Show LAG Statistics

[Table 108](#) describes the output fields for the **show lag statistics** command.

Table 108 Output Fields: Show Lag Statistics

Label	Description
LAG ID	The LAG or multi-link trunk (MLT) that the port is assigned to.
Port ID	The port ID.
Input Bytes	The number of incoming bytes for the LAG on a per-port basis.

Table 108 Output Fields: Show Lag Statistics (Continued)

Label	Description (Continued)
Input Packets	The number of incoming packets for the LAG on a per-port basis.
Output Bytes	The number of outbound bytes for the LAG on a per-port basis.
Output Packets	The number of outbound packets for the LAG on a per-port basis.
Input/Output Errors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character- oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol. For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors.
Totals	Displays the column totals for bytes, packets, and errors.

Sample Output: Show LAG <lag-id> Associations

```
A:ALA-1# show lag 5 associations
=====
Interface Table
=====
Router/ServiceId          Name                      Encap Val
-----
Router: Base              LAG2West                 0
-----
Interfaces
=====
A:ALA-1#
```

Sample Output: Show LAG <lag-id> detail (with MC-LAG Output)

```
*A:pc5# show lag 2 detail
=====
LAG Details
=====
Description:
-----
Details
-----
Lag-id      : 2                Mode           : access
Adm         : up              Opr            : up
Thres. Exceeded Cnt : 2        Port Threshold : 0
Thres. Last Cleared : 04/11/2007 21:50:55 Threshold Action : down
Dynamic Cost  : false         Encap Type     : dot1q
Configured Address : 8e:8b:ff:00:01:42 Lag-IfIndex    :
1342177282
Hardware Address : 8e:8b:ff:00:01:42 Adapt Qos     :
distribute
Hold-time Down : 0.0 sec
LACP          : enabled      Mode           : active
```

```

LACP Transmit Intvl : fast          LACP xmit stdby   : enabled
Selection Criteria  : highest-count  Slave-to-partner  : disabled
Number of sub-groups: 2             Forced            : -
System Id           : 8e:8b:ff:00:00:00 System Priority   : 32768
Admin Key           : 32768          Oper Key           : 32768
Prtr System Id      : 8e:89:ff:00:00:00 Prtr System Priority : 32768
Prtr Oper Key       : 32768

MC Peer Address     : 10.10.10.101    MC Peer Lag-id     : 2
MC System Id        : 01:01:01:01:01:01 MC System Priority  : 2
MC Admin Key        : 1               MC Active/Standby  : active
MC Lacp ID in use   : false           MC extended timeout : false
MC Selection Logic   : waiting for peer info MC Config Mismatch : no mismatch
-----
Port-id      Adm   Act/Stdby Opr   Primary  Sub-group  Forced
Prio
-----
1/1/1        up    active   up    yes      7          -      99
1/1/2        up    standby  down  yes      8          -     100
-----
Port-id      Role    Exp  Def  Dist  Col  Syn  Aggr  Timeout
Activity
-----
1/1/1        actor   No   No   Yes   Yes   Yes   Yes   Yes   Yes
1/1/1        partner No   No   Yes   Yes   Yes   Yes   Yes   Yes
1/1/2        actor   No   No   No    No    No    Yes   Yes   Yes
1/1/2        partner No   No   No    No    Yes   Yes   Yes   Yes
=====
*A:pc5#

```

Sample Output: show lag <lag-id> detail (Without MC-LAG Output)

```

*A:pc5# show lag 2 detail
=====
LAG Details
=====
Description:
-----
Details
-----
Lag-id       : 2          Mode           : access
Adm          : up        Opr            : up
Thres. Exceeded Cnt : 4    Port Threshold : 0
Thres. Last Cleared : 04/11/2007 02:03:49 Threshold Action : down
Dynamic Cost    : false   Encap Type     : dot1q
Configured Address : 8e:8b:ff:00:01:42 Lag-IfIndex    :
1342177282
Hardware Address : 8e:8b:ff:00:01:42 Adapt Qos     :
distribute
Hold-time Down  : 0.0 sec
LACP           : enabled  Mode           : active
LACP Transmit Intvl : fast  LACP xmit stdby : enabled
Selection Criteria : highest-count Slave-to-partner : disabled
Number of sub-groups: 2    Forced            : -
System Id       : 8e:8b:ff:00:00:00 System Priority  : 32768
Admin Key       : 32768    Oper Key        : 32768
Prtr System Id  : 8e:89:ff:00:00:00 Prtr System Priority : 32768
Prtr Oper Key   : 32768

```

```

-----
Port-id      Adm    Act/Stdb  Opr    Primary  Sub-group  Forced
Prio
-----
1/1/1        up     active   up     yes      7          -      99
1/1/2        up     standby  down                    8          -      100
-----

Port-id      Role      Exp    Def    Dist    Col    Syn    Aggr    Timeout
Activity
-----
1/1/1        actor     No     No     Yes     Yes     Yes     Yes     Yes     Yes
1/1/1        partner   No     No     Yes     Yes     Yes     Yes     Yes     Yes
1/1/2        actor     No     No     No      No      No      Yes     Yes     Yes
1/1/2        partner   No     No     No      No      Yes     Yes     Yes     Yes
=====
*A:pc5#

*A:Dut-A# show lag 2 associations per-link-hash sap
=====
SAP Associations
=====
SvcId      SAP                      Active Link                      Oper  Oper
Class      Weight
-----
2          lag-2:4                  1/1/1                          1     500
2          lag-2:5                  1/1/1                          1     100
2          lag-2:6                  1/1/26                         1    1000
2          lag-2:7                  1/1/25                         1    1000
=====
Number of SAP associations: 4

A:bksim4001# show lag 1 per-link-hash
=====
Per-link-hash Weight
=====
Port              Class    Num Users  Agg Weight
-----
1/1/1              1        0          0
1/1/1              2        0          0
1/1/1              3        0          0
=====
Number of entries: 3
=====

```

Sample Output: show lag <lag-id> lcap-partner

```

A:ALU-Dut1# show lag 3 lacp-partner
=====
LAG Partner information
=====
Partner system ID      : ea:3e:ff:00:00:00
Partner system priority : 32768
Partner operational key  : 2
=====

LAG 3 Ports Partner operational information
=====
Port              Actor Port  Prio  Key

```

```

                                port
-----
1/1/52                        33908 33909 5      2
1/1/54                        33910 33911 5      2
1/1/56                        33912 33913 7      2
=====
LAG 3 Ports Partner operational state information
=====
Port                          Exp  Def  Dist Col  Syn  Aggr Time Act
                                out
-----
1/1/52                        No   No   Yes  Yes  Yes  Yes  Yes  Yes
1/1/54                        No   No   Yes  Yes  Yes  Yes  Yes  Yes
1/1/56                        No   No   No   No   No   Yes  Yes  Yes
=====
A:ALU-Dut1#

```

```

A:Dut-A# show lag 10 lacp-neighbors
=====
LAG Neighbor information
=====
Partner system ID             : de:41:ff:00:00:00
Partner system priority       : 32768
Partner operational key       : 32768
=====
LAG port 1/1/6 partner information
-----
Actor port                    : 33862
Partner admin system prio     : 32768
Partner oper system prio      : 32768
Partner admin system ID       : 00:00:00:00:00:00
Partner oper system ID        : de:41:ff:00:00:00
Partner admin key              : 0
Partner oper key              : 32768
Partner admin port            : (Not Specified)
Partner oper port             : 33863
Partner admin port prio       : 32768
Partner oper port prio        : 32768
Partner admin state           : (Not Specified)
Partner oper state            : lacp-timeout aggregation synchronization
                                collecting distributing
=====
A:Dut-A#

```

```

*A:bksim4001>config>lag# selection-criteria highest-weight subgroup-hold-
time 1show lag 1 detail
ght subgroup-hold-time 10
=====
LAG Details
=====
Description                   : To Sim4002
-----
Details
-----
Lag-id                        : 1                               Mode                : access

```

```

Adm                : down                      Opr                : down
Thres. Exceeded Cnt : 0                        Port Threshold       : 0
Thres. Last Cleared  : 01/21/2014 09:00:48      Threshold Action     : down
Dynamic Cost        : false                     Encap Type           : null
Configured Address  : 36:95:ff:00:01:41         Lag-IfIndex          : 1342177281
Hardware Address    : 36:95:ff:00:01:41         Adapt Qos (access)  : distribute
Hold-time Down     : 0.0 sec                    Port Type            : standard
Per-Link-Hash       : disabled
Include-Egr-Hash-Cfg : enabled
Per FP Ing Queuing  : disabled                  Per FP Egr Queuing   : disabled
Per FP SAP Instance : disabled
LACP                : enabled                    Mode                 : passive
LACP Transmit Intvl : fast                      LACP xmit stdby      : enabled
Selection Criteria   : highest-weight           Slave-to-partner     : disabled
Subgrp hold time    : 20.0 sec                  Remaining time       : 2.6 sec
Subgrp selected     : 1                         Subgrp candidate     : 2
Subgrp count        : 2                         Forced               : -
System Id           : 36:95:ff:00:00:00         System Priority      : 32768
Admin Key           : 32768                      Oper Key             : 32768
Prtr System Id      :                          Prtr System Priority : 0
Prtr Oper Key       : 0
Standby Signaling   : lacp
Port weight (gbps)  : (Not Specified)
Weight Threshold    : 0                         Threshold Action     : down
...
=====

```

*A:Dut-A# show lag 2 associations per-link-hash sap

=====

SvcId	SAP	Active Link	Oper Class	Oper Weight
2	lag-2:4	1/1/1	1	500
2	lag-2:5	1/1/1	1	100
2	lag-2:6	1/1/26	1	1000
2	lag-2:7	1/1/25	1	1000

=====

Number of SAP associations: 4

A:bksim4001# show lag 1 per-link-hash

=====

Port	Class	Num Users	Agg Weight
1/1/1	1	10	10
1/1/1	2	0	0
1/1/1	3	2	500

=====

Number of entries: 3

Output Fields: Show LAG <lag-id> Associations

[Table 109](#) describes the output fields for the **show lag <lag-id> associations** command.

Table 109 Output Fields: Show Lag <lag-id> Associations

Label	Description
Service ID	The service associated with the LAG.
Name	The name of the IP interface.
Encap Val	The Dot1q or QinQ values of the port for the IP interface.

2.21.2.7 Licensing Show Command

licensing

Syntax **licensing entitlements**
licensing *port-id*

Context show

Description This command displays information related to the active application licenses on the system.

When the **entitlements** keyword is provided, the display shows all the entitlements in the active licenses in the system, the type of license, the quantity of POOL licenses and the number of each that is in use or available for use.

Parameters **entitlements** — Specifies the number of instances allowed by the active license.
port-id — The port identifier in the format *slot[/mda[/connector[/port]]]*.

<i>slot</i>	1 to 10
<i>mda</i>	1 to 2
<i>connector</i>	c1 to c36
<i>port</i>	1 to 10

Output The following outputs are examples of licensing information.

Sample Output: show licensing entitlements

```
*A:bkvm20# show licensing entitlements
```

```
=====
License                               Available    In-Use      State
-----
MDA Upgrades
  cr1200g-cr1600g                      1           0          VALID
  cr1200g-er1200g                      2           0          VALID
```

```

cr1600g-cr2400g          1          0      VALID
cr1600g-er1600g          0          2      VALID
cr2400g-er2400g          1          0      VALID
er1200g-er1600g          1          0      VALID
er1200g-he1200g          1          0      VALID
er1600g-er2400g          4          0      VALID
er1600g-he1600g          1          0      VALID
er2400g-he2400g          1          0      VALID
he1200g-he1600g          1          0      VALID
he1600g-he2400g          1          0      VALID
=====
*

```

Sample output: show licensing

```

*A:Dut-C# show licensing 2/1
=====
Connector          MAC  Licensed  Restrictions
-----
2/1/c1              1    Yes      400 Gbps per MAC Chip
2/1/c2              1    Yes      400 Gbps per MAC Chip
2/1/c3              1    Yes      400 Gbps per MAC Chip
2/1/c4              2    Yes      400 Gbps per MAC Chip
2/1/c5              2    Yes      400 Gbps per MAC Chip
2/1/c6              2    Yes      400 Gbps per MAC Chip
2/1/c7              3    Yes      400 Gbps per MAC Chip
2/1/c8              3    Yes      400 Gbps per MAC Chip
2/1/c9              3    Yes      400 Gbps per MAC Chip
2/1/c10             4    Yes      400 Gbps per MAC Chip
2/1/c11             4    Yes      400 Gbps per MAC Chip
2/1/c12             4    Yes      400 Gbps per MAC Chip
2/1/c13             5    Yes      400 Gbps per MAC Chip
2/1/c14             5    Yes      400 Gbps per MAC Chip
2/1/c15             5    Yes      400 Gbps per MAC Chip
2/1/c16             6    Yes      400 Gbps per MAC Chip
2/1/c17             6    Yes      400 Gbps per MAC Chip
2/1/c18             6    Yes      400 Gbps per MAC Chip
2/1/c19             7    Yes      400 Gbps per MAC Chip
2/1/c20             7    Yes      400 Gbps per MAC Chip
2/1/c21             7    Yes      400 Gbps per MAC Chip
2/1/c22             8    Yes      400 Gbps per MAC Chip
2/1/c23             8    Yes      400 Gbps per MAC Chip
2/1/c24             8    Yes      400 Gbps per MAC Chip
2/1/c25             9    Yes      400 Gbps per MAC Chip
2/1/c26             9    Yes      400 Gbps per MAC Chip
2/1/c27             9    Yes      400 Gbps per MAC Chip
2/1/c28             10   Yes      400 Gbps per MAC Chip
2/1/c29             10   Yes      400 Gbps per MAC Chip
2/1/c30             10   Yes      400 Gbps per MAC Chip
2/1/c31             11   Yes      400 Gbps per MAC Chip
2/1/c32             11   Yes      400 Gbps per MAC Chip
2/1/c33             11   Yes      400 Gbps per MAC Chip

```

Output Fields

[Table 110](#) describes the output fields for the **show licensing** command.

Table 110 **Output Fields: Show Licensing**

Label	Description
License	The description of the license.
In-Use	The number of instances that are currently being used.
Available	The number of instances that are not currently being used. "License err" if in use is greater than the number of the allowed entitlements.
State	Indicates the state of the license as: • VALID - license can be used in the node. • VIOLATION - license is in use but not covered by the currently activated license key file.
Connector	The identity of connector for the license limits.
MAC	The index of the MAC on the card or MDA to which the connector is attached.
Licensed	Whether or not the connector is licensed. Yes — the connector is included in the licensed level of the card or MDA and may be configured for a breakout. Other rules may apply as described in the restrictions field. No — this connector is not included in the licensed level of the card or MDA and it cannot be configured for breakout.
Restrictions	Additional license rules that may apply to the use of the connector, for example, bandwidth limits to the set of connectors on the same MAC.

2.21.2.8 MACsec Show Commands

macsec

Syntax **macsec**

Context show

Description This command displays MACsec information.

connectivity-association

- Syntax** **connectivity-association** [*ca-name*] [**detail**]
- Context** show>macsec
- Description** This command displays MACsec specific CA information.
- Parameters** *ca-name* — Displays CA name information up to 256 characters in length.
detail — Displays MACsec CA detailed information.
- Output** See the following sections for output samples:
- [Sample Output: show macsec connectivity-association](#)
 - [Output Fields: show macsec connectivity-association](#)
 - [Sample Output: show macsec connectivity-association <ca-name>](#)
 - [Output Fields: show macsec connectivity-association <ca-name>](#)
 - [Sample Output: show macsec connectivity-association <ca-name> detail](#)

Sample Output: show macsec connectivity-association

```
A:Dut-C# show macsec connectivity-association
=====
ca-name           : dut_B_C_128_01
ca-name           : dut_B_C_256_01
ca-name           : dut_B_C_128_xpn_01
ca-name           : dut_B_C_256_xpn_01
ca-name           : to_Juniper_1_1_2__1
ca-name           : abcdefghijklmnoprstuvwxyz@!
=====
```

Output Fields: show macsec connectivity-association

[Table 111](#) describes the output fields for the **show macsec connectivity-association** command.

Table 111 Output Fields: show macsec connectivity-association

Label	Description
ca-name	Specifies the CA name.

Sample Output: show macsec connectivity-association <ca-name>

```
A:Dut-C# show macsec connectivity-association "abcdefghijklmnoprstuvwxyz@!"
=====
Connectivity Association "abcdefghijklmnoprstuvwxyz@!"
=====
Admin State       : Up
Description       : alsfjalsfjafja;lsjflasjflasjfl
```

```

Replay Protection : Disabled
Replay Window Size : 333
Macsec Encrypt : Enabled
Clear Tag Mode : dual-tag
Cipher Suite : gcm-aes-256
Encryption Offset : 30
Assigned ports : 2/1/9 2/1/10

```

```
-----
Static Cak
-----
```

```

MKA Key Server Priority : 16
Active Pre-Shared-Key Index : 1
Active Pre-Shared-Key CKN : aabbccddeeff00112233445566778899
=====

```

Output Fields: show macsec connectivity-association <ca-name>

[Table 112](#) describes the output fields for the **show macsec connectivity-association <ca-name>** command.

Table 112 Output Fields: show macsec connectivity-association <ca-name>

Label	Description
Admin State	Up — The CA is administratively up. Down — The CA is administratively down. If port <x/y/z> ethernet>macsec is shutdown, the admin state will be down. Otherwise, the admin state is up.
Description	Specifies a user description for this CA.
Replay Protection	Enabled — Replay Protection is enabled. Disabled — Replay Protection is disabled. If replay protection is enabled for this CA, the out of window packet will be discarded.
Replay Window Size	Specifies the size, in packets, of the replay window.
Macsec Encrypt	Enabled — MACsec is enabled. Disabled — MACsec is disabled.
Clear Tag Mode	Specifies the clear tag mode: single-tag, dual-tag.
Cipher Suite	Specifies the cipher suite used for encrypting the SAK: gcm-aes-128, gcm-aes-256, gcm-aes-xpn-128, gcm-aes-xpn-256.
Encryption Offset	Specifies the encryption offset configured on this node: 0, 30, 50.
Assigned ports	Specifies all ports that contain this CA.

Table 112 Output Fields: show macsec connectivity-association (Continued)<ca-name> (Continued)

Label	Description
MKA Key Server Priority	Specifies the MKA key server priority: 0-255 (default 16).
Active Pre-Shared Key Index	Specifies the active pre-shared key index: 1-2 (default 1).
Active Pre-Shared Key CKN	Specifies the active PSK CAK name.

Sample Output: show macsec connectivity-association <ca-name> detail

```
A:Dut-C# show macsec connectivity-association "abcdefghijklmnoprstuvwxyz@!" detail
=====
Connectivity Association "abcdefghijklmnoprstuvwxyz@!"
=====
Admin State          : Up
Description          : alsfjalsfjafja;lsjflasjflasjfl
Replay Protection    : Disabled
Replay Window Size   : 333
Macsec Encrypt       : Enabled
Clear Tag Mode       : dual-tag
Cipher Suite         : gcm-aes-256
Encryption Offset    : 30
Assigned ports       : 2/1/9 2/1/10
-----
Static Cak
-----
MKA Key Server Priority      : 16
Active Pre-Shared-Key Index : 1
Active Pre-Shared-Key CKN   : aabbccddeeff00112233445566778899
=====
```

mka-session

- Syntax** **mka-session** [**port** *port-id*]
mka-session [**port** *port-id*] **detail**
mka-session [**port** *port-id*] **statistics**
- Context** show>macsec
- Description** This command displays MACsec MKA session information.
- Parameters** *port-id* — Specifies the port ID, up to 17characters.
detail — Displays MACsec MKA session detailed information.
statistics — Displays MACsec MKA session statistical information.

Output See the following sections for output samples:

- [Sample Output: show macsec mka-session port](#)
- [Output Fields: show macsec mka-session port](#)
- [Sample Output: show macsec mka-session port detail](#)
- [Output Fields: show macsec mka-session port detail](#)
- [Sample Output: show macsec mka-session detail](#)
- [Sample Output: show macsec mka-session statistics](#)

Sample Output: show macsec mka-session port

```
A:Dut-C# show macsec mka-session port 2/1/11
=====
MKA Session for port 2/1/11
=====
Port           : 2/1/11
Security Zone   : 3
=====
Live Peer List
=====
Member Identifier      Mesg Num  Rx-SCI              KS priority
-----
bf4102704294fa1057022bdf    28322    a47b2ce112ef0000      16
=====
Potential Peer List
=====
Member Identifier      Mesg Num  Rx-SCI              KS priority
-----
```

Output Fields: show macsec mka-session port

[Table 113](#) describes the output fields for the following **show macsec mka-session port** command.

Table 113 Output Fields: show macsec mka-session port

Label	Description
MKA Session for port	Display the MKA session for the current port
Port	Specifies the MKA session current port.
Security Zone	Specifies Security Zone does this port belongs to.
Live Peer List	Specifies Peers (Participants) which have provided their MI and MN via KMA. The peer entry is in the Live Peer List.
Member Identifier	Specifies the MI of the peer entry.

Table 113 Output Fields: show macsec mka-session port (Continued)

Label	Description
Mesg Num	Specifies the latest Member Number of the peer entry.
Rx-SCI	Specifies the Peer Rx-SCI.
KS-priority	Specifies the Peer Key server priority.
Potential Peer List	Peers (Participants) which have Potential Peers List includes all the other peers that have transmitted an MKPDU that has been directly received by the participant or that were included in the Live Peers List of a MKPDU transmitted by a peer that has proved liveness, an MKA PDU. The peer entry is in the Potential Peer List.

Sample Output: show macsec mka-session port detail

```
A:Dut-C# show macsec mka-session port 2/1/11 detail
=====
MKA Session for port 2/1/11
=====
Port                : 2/1/11
Security Zone       : 3
MKA Oper State      : unknown value
Oper Cipher Suite   : unknown value
Oper Encrypt Offset : 0
CAK Name            : 11223344556677889900aabbccddeeff11223344556677889900aabbcc*
MKA Member ID       : f134218784b114eb61dbe834
Transmit Interval   : 2000
Outbound SCI        : a4:7b:2c:e1:12:8f
Message Number      : 28298
Key Number          : 878
Key Server          : yes
Key Server Priority  : 16
Latest SAK AN       : 3
Latest SAK KI       : f134218784b114eb61dbe8340000036d
Previous SAK AN     : 2
Previous SAK KI     : f134218784b114eb61dbe83400000000
=====
* indicates that the corresponding row element may have been truncated.
=====
Live Peer List
=====
Member Identifier      Mesg Num  Rx-SCI                KS priority
-----
bf4102704294fa1057022bdf  28323    a47b2ce112ef0000      16
=====
Potential Peer List
=====
Member Identifier      Mesg Num  Rx-SCI                KS priority
-----
=====
MKA Session Statistics for port 2/1/11
=====
```

```

Peer Removed Due to Timeout : 0
CKN Not Found                : 0
New Live peer                : 0
SAK Generated by Server      : 0
SAK Installed for TX         : 0
SAK Installed for RX         : 0
PDU Too Small                : 0
PDU Too Big                  : 0
PDU Not Quad Size            : 0
PDU Message Number Invalid   : 0
PDU Param Set Size Invalid   : 0
PDU Liveness Check Fail      : 0
Param Set Not Quad Size      : 0
Unsupported Agility          : 0
Invalid CAK Name Length      : 0
ICV Check Failed             : 0
Peer Using Same MID          : 0
SAK From Non-Live Peer       : 0
SAK From Non-Key Server      : 0
SAK Decrypt Fail             : 0
SAK Encrypt Fail             : 0
Key Number Invalid           : 0
SAK Installation Failed       : 0
CAK Info Missing             : 0
Max Peers Set as Zero        : 0
=====

```

Output Fields: show macsec mka-session port detail

Table 114 describes the output fields for the following commands:

- **show macsec mka-session port**
- **show macsec mka-session detail**
- **show macsec mka-session statistics**

Table 114 **Output Fields: show macsec mka-session port detail**

Label	Description
MKA Oper State	Specifies the operational state of the MKA participant on this port. The operational MKA state will be up if MKA hellos are received on this port and have a valid session.
Oper Cipher Suite	Specifies the operational encryption algorithm used for datapath PDUs when all parties in the CA have the (SAK). This value is specified by the key server: gcm-aes-128, gcm-aes-256, gcm-aes-xpn-128, gcm-aes-xpn-256.
Oper Encrypt Offset	Specifies the operational encryption offset used for the datapath PDUs when all parties in the CA have the SAK. This value is specified by the key server: 0, 30, 50.
CAK Name	Specifies the name of the CAK in use by this MKA which is used to find the correct CAK.

Table 114 **Output Fields: show macsec mka-session port detail**

Label	Description
MKA Member ID	Specifies indicates the Member Identifier (MI) for the MKA instance.
Transmit Interval	Specifies the time interval (in ms) at which the MKA broadcasts its liveliness to its peers and is non-configurable.
Outbound SCI	Specifies the Secure Channel Identifier (SCI) information for transmitting MACsec frames and consists of the outgoing port MAC Address and a port identifier.
Message Number	Specifies the current count of MKA messages that is attached to MKA PDUs.
Key Number	Specifies the number of the currently assigned CAK. When a new CAK is generated, this number is incremented. A SAK is identified by 128-bit Key Identifier (KI) and 32-bit Key-Number (KN).
Key Server	Specifies whether this server is the highest priority server in the peer group: no, yes.
Key Server Priority	Specifies the priority of the active key server: 0-255 (default 16).
Latest SAK AN	Specifies the Association Number (AN) of the latest Secure Association Key (SAK). This number is concatenated with an SCI to identify a Secure Association (SA). In SR OS, only 2 SAKs are supported.
Latest SAK KI	Specifies the Key Identifier (KI) of the latest SAK. This number is derived from the MI of the key server and the key number.
Previous SAK AN	Specifies the AN of the previous SAK. This number is concatenated with an SCI to identify an SA.
Previous SAK KI	Specifies the KI of the previous SAK. This number is derived from the MI of the key server and the key number.
Peer Removed Due to Timeout	Specifies the number of peers removed from the live/potential peer list due to not receiving an MKPDU within the MKA Live Time (6.0 seconds) and is not configurable.
CKN Not Found	Specifies the number of MKPDUs received with a CKN that does not match the CA configured for the port.
New Live Peer	Specifies the number of validated peers that have been added to the live peer list.
SAK Generated by Server	Specifies the number of SAKs generated by this MKA instance.
SAK Installed for TX	Specifies the number of SAKs installed for transmitting.

Table 114 **Output Fields: show macsec mka-session port detail**

Label	Description
SAK Installed for RX	Specifies the number of SAKs installed for receiving.
PDU Too small	Indicates that the number of MKPDUs received that are less than 32 octets.
PDU Too big	Indicates the number of MKPDUs received where the EAPOL header indicates a size larger than the received packet.
PDU Not Quad Size	Indicates the number of MKPDUs received with a size that is not a multiple of 4 octets long.
PDU Message Number Invalid	Indicates the number of MKPDUs received out of order as indicated by the Message Number.
PDU Param Set Size Invalid	Indicates the number of MKPDUs received which contain a parameter set body length that exceeds the remaining length of the MKPDU.
PDU Liveness Check Fail	Indicates the number of MKPDUs received which contain an MN that is not acceptably recent.
Param Set Not Quad Size	Indicates the number of MKPDUs received which contain a parameter set that is not a multiple of 4 octets long.
Unsupported Agility	Indicates the number of MKPDUs received which contain an unsupported Algorithm Agility value.
Invalid CAK Name Length	Indicates the number of MKPDUs received which contain a CAK name that exceeds the maximum CAK name length.
ICV Check Failed	Indicates the number of MKPDUs received which contain an ICV value that does not authenticate.
Peer Using Same MID	Indicates the number of MKPDUs received which contain a peer list with an MI entry which conflicts with the local MI.
SAK From Non-Live Peer	Indicates the number of SAKs received from peer that is not a member of the Live Peers List.
SAK From Non-Key Server	Indicates the number of SAKs received from an MKA participant that has not been designated as the Key Server. Only the key server should distribute SAK.
SAK Decrypt Fail	Indicates the number of AES Key Wrap SAK decryption failures that have occurred.
SAK Encrypt Fail	Indicates the number of AES Key Wrap SAK encryption failures that have occurred.
Key Number Invalid	Indicates the number of SAKs received with an invalid Key Number.

Table 114 Output Fields: show macsec mka-session port detail

Label	Description
SAK Installation Failed	Indicates the number of Secy SAK installation failures that have occurred.
CAK Info Missing	Indicates the number of times internal CAK data is not available for the generation of the SAK.
Max Peers Set as Zero	Indicates the number of Secy SAK installations that have failed due to the max peer entry being set to 0.

Sample Output: show macsec mka-session detail

```
A:Dut-C# show macsec mka-session detail
=====
MKA Session for port 2/1/11
=====
Port                : 2/1/11
Security Zone       : 3
MKA Oper State      : unknown value
Oper Cipher Suite   : unknown value
Oper Encrypt Offset : 0
CAK Name            : 11223344556677889900aabbccddeeff11223344556677889900aabbcc*
MKA Member ID       : f134218784b114eb61dbe834
Transmit Interval   : 2000
Outbound SCI        : a4:7b:2c:e1:12:8f
Message Number      : 28285
Key Number          : 878
Key Server          : yes
Key Server Priority  : 16
Latest SAK AN       : 3
Latest SAK KI       : f134218784b114eb61dbe83400000036d
Previous SAK AN      : 2
Previous SAK KI      : f134218784b114eb61dbe83400000000
=====
* indicates that the corresponding row element may have been truncated.
=====
Live Peer List
=====
Member Identifier      Mesg Num  Rx-SCI                KS priority
-----
bf4102704294fa1057022bdf    28310    a47b2ce112ef0000      16
=====
Potential Peer List
=====
Member Identifier      Mesg Num  Rx-SCI                KS priority
-----
=====
MKA Session Statistics for port 2/1/11
=====
Peer Removed Due to Timeout : 0
CKN Not Found               : 0
New Live peer                : 0
```

```

SAK Generated by Server      : 0
SAK Installed for TX         : 0
SAK Installed for RX         : 0
PDU Too Small                : 0
PDU Too Big                  : 0
PDU Not Quad Size            : 0
PDU Message Number Invalid   : 0
PDU Param Set Size Invalid   : 0
PDU Liveness Check Fail      : 0
Param Set Not Quad Size      : 0
Unsupported Agility          : 0
Invalid CAK Name Length      : 0
ICV Check Failed             : 0
Peer Using Same MID          : 0
SAK From Non-Live Peer       : 0
SAK From Non-Key Server      : 0
SAK Decrypt Fail             : 0
SAK Encrypt Fail             : 0
Key Number Invalid           : 0
SAK Installation Failed      : 0
CAK Info Missing             : 0
Max Peers Set as Zero        : 0

```

```
=====
```

Sample Output: show macsec mka-session statistics

```
A:Dut-C# show macsec mka-session statistics
```

```
=====
MKA Session Statistics for port 2/1/11
```

```

=====
Peer Removed Due to Timeout : 0
CKN Not Found               : 0
New Live peer               : 0
SAK Generated by Server     : 0
SAK Installed for TX        : 0
SAK Installed for RX        : 0
PDU Too Small               : 0
PDU Too Big                 : 0
PDU Not Quad Size           : 0
PDU Message Number Invalid  : 0
PDU Param Set Size Invalid  : 0
PDU Liveness Check Fail     : 0
Param Set Not Quad Size     : 0
Unsupported Agility          : 0
Invalid CAK Name Length     : 0
ICV Check Failed            : 0
Peer Using Same MID         : 0
SAK From Non-Live Peer      : 0
SAK From Non-Key Server     : 0
SAK Decrypt Fail            : 0
SAK Encrypt Fail            : 0
Key Number Invalid          : 0
SAK Installation Failed     : 0
CAK Info Missing            : 0
Max Peers Set as Zero       : 0
=====

```

2.21.2.9 Monitor Commands

card

Syntax	card <i>slot-number</i> fp <i>fp-number</i> ingress { access network } queue-group <i>queue-group-name</i> instance <i>instance-id</i> [interval <i>seconds</i>] [repeat <i>repeat</i>] policer <i>policer-id</i> [absolute percent-rate [<i>reference-rate</i>]]
Context	monitor
Description	This command monitors card parameters.
Parameters	<i>slot-number</i> — Specifies the slot number associated with the queue group, expressed as an integer. Values 1 to 10 <i>fp-number</i> — Specifies the FP number associated with the queue group, expressed as an integer. Values 1 to 8 ingress — Displays policer statistics applied on the ingress FP. access — Displays policer statistics on the FP access. network — Displays policer statistics on the FP network. <i>queue-group-name</i> — Specifies the queue group name up to 32 characters. <i>instance-id</i> — Specifies the identification of a specific instance of the queue group. Values 1 to 65535 <i>seconds</i> — Configures the interval for each display in seconds. Values 11 to 60 Default 11 <i>repeat</i> — Configures how many times the command is repeated. Values 1 to 999 Default 10 <i>policer-id</i> — Must exist within the queue-group template applied to the ingress context of the forwarding plane. Values 1 to 32 absolute — Specifies the raw statistics to display, without processing. No calculations are performed on the delta or rate statistics. percent-rate — Specifies the rate-per-second for each statistic is displayed based on the reference rate of 10G.

reference-rate — Rate-per-second for each statistic is displayed as a percentage based on the reference rate specified.

Values 100M, 1G, 10G, 40G, 100G, 400G
Default 10G

port

Syntax **port** *port-id* [*port-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] [**multiclass**]

Context monitor

Description This command enables port traffic monitoring. The specified port(s) statistical information displays at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the specified port(s). The subsequent statistical information listed for each interval is displayed as a delta to the previous display.

When the **rate** is specified, the "rate per second" for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands but only statistical information displays. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters *port-id* — Specify up to 5 port IDs. Port-IDs are only MLPPP bundles or bundle protection groups when the multiclass is specified.

<i>port-id</i>	<i>slot/mda/port</i> [<i>.channel</i>]	
	eth-sat-id	esat-id/slot/port
		esat
		id
		keyword
		1 to 20
	pxc-id	pxc-id.sub-port
		pxc
		keyword
		id
		1 to 64
		sub-port
		a, b
	aps-id	aps-group-id[.channel]
		aps
		keyword
		group-id
		1 to 64 (16 for 7750 SR-c12 and 7750 SR-c4)
	bundle ID	bundle-type-slot/mda.bundle-num
		bundle
		keyword
		type
		ima, ppp
		bundle-num
		1 to 128 (16 for 7750 SR-c12 and 7750 SR-c4)

bpgrp-type-bpgrp-num

seconds — Configures the interval for each display in seconds.

Values 3 to 60

Default 10 seconds

repeat — Configures how many times the command is repeated.

Values 1 to 999

Default 10

absolute — When the **absolute** is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate — When the **rate** is specified, the rate-per-second for each statistic is displayed instead of the delta.

Output The following output is an example of port multiclass information.

Sample Output

```
A:ALA-12>monitor# port 2/1/4 interval 3 repeat 3 absolute
=====
Monitor statistics for Port 2/1/4
=====
                                     Input                               Output
-----
At time t = 0 sec (Base Statistics)
-----
Octets                               0                               0
Packets                             39                              175
Errors                               0                               0
-----
At time t = 3 sec (Mode: Absolute)
-----
Octets                               0                               0
Packets                             39                              175
Errors                               0                               0
-----
At time t = 6 sec (Mode: Absolute)
-----
Octets                               0                               0
Packets                             39                              175
Errors                               0                               0
-----
At time t = 9 sec (Mode: Absolute)
-----
Octets                               0                               0
Packets                             39                              175
Errors                               0                               0
=====
A:ALA-12>monitor#

A:ALA-12>monitor# port 2/1/4 interval 3 repeat 3 rate
```

```
=====
Monitor statistics for Port 2/1/4
=====
                               Input                               Output
-----
At time t = 0 sec (Base Statistics)
-----
Octets                        0                                0
Packets                       39                               175
Errors                        0                                0
-----
At time t = 3 sec (Mode: Rate)
-----
Octets                        0                                0
Packets                       0                                0
Errors                        0                                0
Bits                          0                                0
Utilization (% of port capacity) 0.00                        0.00
-----
At time t = 6 sec (Mode: Rate)
-----
Octets                        0                                0
Packets                       0                                0
Errors                        0                                0
Bits                          0                                0
Utilization (% of port capacity) 0.00                        0.00
-----
At time t = 9 sec (Mode: Rate)
-----
Octets                        0                                0
Packets                       0                                0
Errors                        0                                0
Bits                          0                                0
Utilization (% of port capacity) 0.00                        0.00
=====
A:ALA-12>monitor#
```

```
=====
*A:Cpm-A> monitor port bundle-fr-1/1.1
=====
Monitor statistics for Port bundle-fr-1/1.1
=====
                               Input                               Output
-----
At time t = 0 sec (Base Statistics)
-----
Octets                        0                                0
Packets                       0                                0
Errors                        0                                0
```

port

- Syntax** port all-ethernet-rates [interval seconds] [repeat repeat]
- Context** monitor

Description This command enables port traffic monitoring and utilization output for all data Ethernet ports enabled in the system, in a table output format. The specified port(s) statistical information displays at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the specified port(s). The subsequent statistical information listed for each interval is displayed as a delta to the previous display.

Monitor commands are similar to show commands but only statistical information displays. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters *seconds* — Configures the interval for each display in seconds.

Values 3 to 60

Default 10 seconds

repeat — Configures how many times the command is repeated.

Values 1 to 999

Default 10

Output The following output is an example of port Ethernet information.

Sample Output

```
A:ALA-12>monitor# port all-ethernet-rates interval 3 repeat 100
=====
Monitor statistics for all Ethernet Port Rates
=====
Port-Id          D          Bits   Packets   Errors   Util
-----
At time t = 0 sec (Base Statistics)
-----
5/1/1            I           0        0         0       0.00
                  O           0        0         0       0.00
5/2/1            I           0        0         0       0.00
                  O           0        0         0       0.00
-----
At time t = 3 sec (Mode: Rate)
-----
5/1/1            I           0        0         0       0.00
                  O           0        0         0       0.00
5/2/1            I           0        0         0       0.00
                  O           0        0         0       0.00
-----
At time t = 6 sec (Mode: Rate)
-----
5/1/1            I           0        0         0       0.00
                  O           0        0         0       0.00
5/2/1            I           0        0         0       0.00
```

```

                                0          0          0          0          0.00

-----
At time t = 9 sec (Mode: Rate)
-----
5/1/1      I      4286480384    697669          0      43.98
           I      4286382080    697653          0      43.98

5/2/1      I      4254070784    692394          0      43.64
           I      4253952000    692375          0      43.64

-----
At time t = 12 sec (Mode: Rate)
-----
5/1/1      I      9746288640    1586K          0      99.99
           I      9746216960    1586K          0      99.99

5/2/1      I      9746280448    1586K          0      99.99
           I      9746167808    1586K          0      99.99

=====
A:ALA-12>monitor#

```

port

Syntax **port** *port-id* [*port-id*] [*interval seconds*] [*repeat repeat*] [*absolute* | *rate*] [*multiclass*]
port *port-id* [*port-id*] **atm** [*aal-5* | *cp cp-id* | *ilmi* | **interface-connection** | **pvc** *vpi/vci* [*aal-5* |
oam] | **pvp** *vpi* [*aal-5* | *oam* | **pvt** *vpi1.vpi2*] [*interval seconds*] [*repeat repeat*] [*absolute*
| *rate*]
port all-ethernet-rates [*interval seconds*] [*repeat repeat*]

Context monitor

Description This command enables port traffic monitoring. Various keywords can be used to monitor specific attributes of the ATM port. The specified port(s) statistical information displays at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the specified port(s). The subsequent statistical information listed for each interval is displayed as a delta to the previous display.

When the **rate** is specified, the "rate per second" for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands but only statistical information displays. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters *port-id* — Specifies up to 5 port IDs. Port-IDs are only MLPPP bundles or bundle protection groups when the multiclass is specified.

```

port-id      slot/mda/port [.channel]
               bundle-id      bundle-type-slot/mda.bundle-num
                               bundle
                               keyword

```


	<i>type</i>	ima, fr, ppp
	<i>bundle-num</i>	1 to 336
bgrp-id	<i>bpgrp-type-bpgrp-num</i>	
	bgrp	keyword
	type	ima, ppp
	bgrp-num	1 to 2000
aps-id	<i>aps-group-id[.channel]</i>	
	aps	keyword
	<i>group-id</i>	1 to 64 (16 for 7750 SR-c12 and 7750 SR-c4)

seconds — Configures the interval for each display in seconds.

Values 3 to 60

Default 10 seconds

repeat — Configures how many times the command is repeated.

Values 1 to 999

Default 10

absolute — When the **absolute** is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate — When the **rate** is specified, the rate-per-second for each statistic is displayed instead of the delta.

atm — Monitors ATM statistics.

aal-5 — Monitors the AAL-5 statistics for the port.

cp-id — Monitors ATM connection profile statistics.

Values 1 to 8000

ilmi — Monitors ATM ILMI statistics.

interface-connection — Monitors ATM interface connection statistics.

pvc *vpi/vci* — Monitors ATM PVC statistics.

Values vpi -[0 to 4095] (NNI)
[0 to 255] (UNI)
vci - [1|2|5 to 65535]

pvp *vpi* — Monitors ATM PVP statistics

Values vpi [0 to 4095] (NNI)
[0 to 255] (UNI)

pvt *vpi1.vpi2* — Monitors ATM PVT statistics.

Values [0 to 4095].[0 to 4095] (NNI)

[0 to 255].[0 to 255] (UNI)

oam — Monitors ATM PVC/PVP OAM statistics.

Output The following output is an example of port ATM information.

Sample Output

```
A:ALA-49# monitor port 9/1/1 atm interval 3 repeat 2 absolute
=====
Monitor ATM statistics for Port 9/1/1
=====
                                     Input      Output
-----
At time t = 0 sec (Base Statistics)
-----
Octets                               0          0
Cells                               0          0
Unknown VPI/VCI Cells                0
-----
At time t = 3 sec (Mode: Absolute)
-----
Octets                               0          0
Cells                               0          0
Unknown VPI/VCI Cells                0
-----
At time t = 6 sec (Mode: Absolute)
-----
Octets                               0          0
Cells                               0          0
Unknown VPI/VCI Cells                0
=====
A:ALA-49#
```

queue-group

Syntax **queue-group** *queue-group-name* **egress** *access* [**instance** *instance-id*] [**egress-queue** *egress-queue-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
queue-group *queue-group-name* **ingress** *access* **ingress-queue** *ingress-queue-id* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]
queue-group *queue-group-name* **egress** *network* **instance** *instance-id* [**policer** *policer-id*] [**egress-queue** *egress-queue-id*] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

Context monitor>port

Description This command enables queue-group monitoring for the specified parameters.

Parameters *queue-group-name* — Specifies the name of the queue-group up to 32 characters.
access — Specifies the access type.

Values access

network — Specifies the network type.

Values network

instance-id — Specifies the identification of a specific instance of the queue group.

Values 1 to 65535

policer-id — The specified policer ID must exist within the queue-group template applied to the ingress context of the forwarding plane.

Values 1 to 8

egress-queue-id — Monitors statistics fir this queue.

Values 1 to 8

ingress-queue-id — Monitors statistics for this queue.

Values 1 to 32

seconds — Configures the interval for each display, in seconds.

Values 11 to 60

Default 11

repeat — Configures how many times the command is repeated.

Values 1 to 999

Default 10

absolute — When the absolute keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate — When the rate keyword is specified, the rate-per-second for each statistic is displayed instead of the delta.

vport

Syntax	vport <i>name</i> [interval <i>seconds</i>] [repeat <i>repeat</i>] [absolute rate] vport <i>name</i> [interval <i>seconds</i>] [repeat <i>repeat</i>] monitor-threshold
Context	monitor>port
Description	This command monitors Vport statistics.
Parameters	<i>name</i> — specifies the Vport name up to 32 characters. <i>seconds</i> — Configures the interval for each display, in seconds. Values 11 to 60 <i>repeat</i> — Configures how many times the command is repeated. Values 1 to 999

absolute — When the absolute keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

rate — When the rate keyword is specified, the rate-per-second for each statistic is displayed instead of the delta.

monitor-threshold — Displays the exceed count for the port-scheduler under Vport (if specified) or for a physical port.

2.21.2.10 Clear Commands

card

Syntax

```
card slot-number
card slot-number fp fp-number hs-pool high-water-mark
card slot-number fp fp-number hs-pool high-water-mark mid-pool [value]
card slot-number fp fp-number hs-pool high-water-mark root-pool [value]
card slot-number fp fp-number hs-pool high-water-mark system
card slot-number fp fp-number dist-cpu-protection
card slot-number soft [hard-reset-unsupported-mdas]
card slot-number fp fp-number fwd-engine drop-reason statistics
card slot-number fp fp-number ingress mode {access | network} queue-group group-name instance instance statistics
```

Context clear

Description This command re-initializes the card in the specified slot. A **clear card** command (without the soft) is referred to as a *Hard Reset*. A **clear card x soft** command (with the soft) is referred to as a *Soft Reset*.

Parameters *slot-number* — Clears information for the specified card slot.

Values 1 to 20, A, B, C, D

fp-number — Displays information for the specified forwarding plane.

Values 1 to 8

high-water-mark — Clears the high water marks within the HS pool on the specified card and FP.

mid-pool — Clears the high water marks within the mid pools in the HS pool and on the specified card and FP.

Values 1 to 16

root-pool — Clears the high water marks within the root pools in the HS pool and on the specified card and FP.

Values 1 to 16

system — Clears the high water marks within the in the HS psystem pool and on the specified card and FP.

dist-cpu-protection — Clears the distributed CPU protection information.

soft — Issues a soft reset of the I/O module (IOM).

hard-reset-unsupported-mdas — Allows a soft reset operation when some of the MDA's cannot perform a soft reset. A soft reset is performed on MDAs that support a soft reset and a hard reset is performed on MDAs that do not support soft resets.

fwd-engine — Specifies to clear the forwarding engine information.

drop-reason — Specifies to clear the drop reason information.

statistics — Specifies to clear the statistics.

ingress — Specifies to clear the ingress queue group information.

mode — Specifies the mode in which the card is to operate.

access — Specifies to clear the access mode information.

network — Specifies to clear the network mode information.

group-name — Specifies the group name, to a maximum of 32 characters.

instance — Specifies the instance of the named queue group to be cleared.

Values 1 to 65535

lag

Syntax	lag <i>lag-id</i> statistics lag <i>lag-id</i> bfd-session family { ipv4 ipv6 } lag-port <i>port-id</i>
Context	clear
Description	This command clears statistics for the specified LAG ID.
Parameters	<i>lag-id</i> — Specifies the LAG ID. Values 1 to 800 statistics — Specifies to clear statistics or the BFD session for the specified LAG ID. bfd-session — Specifies to clear a BFD session. family — Specifies the IP address family. ipv4 — Specifies IPv4. ipv6 — Specifies IPv6.

port-id — Specifies the port ID.

Values *slot* [/*mda* [/*port*]]

mda

Syntax	mda <i>mda-id</i> [statistics]
Context	clear
Description	This command reinitializes the specified MDA in a particular slot.
Parameters	<i>mda-id</i> — Clears the specified slot and MDA/CMA.
	Values <i>slot/mda</i>
	slot: 1 to 10
	mda: 1, 2
	statistics — Clears statistics for the specified MDA.

port

Syntax	port { <i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>aps-id</i> } atm port <i>port-id</i> atm cp [<i>cp</i>] statistics port <i>port-id</i> atm ilmi statistics port <i>port-id</i> atm interface-connection statistics port <i>port-id</i> atm pvc [<i>vpi</i> [/ <i>vci</i>]] statistics port <i>port-id</i> atm pvp [<i>vpi</i>] statistics port <i>port-id</i> atm pvt [<i>vpi1.vpi2</i>] statistics port { <i>port-id</i> <i>bundle-id</i> <i>bpgrp-id</i> <i>aps-id</i> } statistics port <i>port-id</i> ethernet dampening port <i>port-id</i> ethernet efm-oam events [local remote] port <i>port-id</i> exp-secondary-shaper <i>shaper-name</i> statistics port <i>port-id</i> fwd-engine drop-reason statistics port <i>port-id</i> hs-secondary-shaper <i>shaper-name</i> statistics port <i>port-id</i> monitor-threshold port <i>port-id</i> hs-pool high-water-mark {[standard <i>value</i>] [alternate <i>value</i>]} port <i>port-id</i> phys-state-change-count port <i>port-id</i> queue-group <i>queue-group-name</i> [instance <i>instance-id</i>] queue-depth [<i>queue</i> <i>queue-id</i>] { ingress egress } [access network] port <i>port-id</i> queue-group <i>queue-group-name</i> [instance <i>instance-id</i>] [access network] { ingress egress } statistics port <i>port-id</i> vport [<i>name</i>] monitor-threshold port <i>port-id</i> vport <i>name</i> statistics

Context	clear																						
Description	This command clears port statistics for the specified ports.																						
Parameters	<i>port-id</i> — Specifies the physical port identifier. Values <i>slot</i> [/mda[/port]] or <i>slot/mda/port</i> [.channel]] <i>bundle-id</i> — Specifies the multilink bundle identifier. Values <table> <tr> <td><i>bundle-type-slot/mda.bundle-num</i></td><td></td></tr> <tr> <td><i>bundle</i></td><td>keyword</td></tr> <tr> <td><i>type</i></td><td>ima, fr, ppp</td></tr> <tr> <td><i>bundle-num</i></td><td>1 to 336</td></tr> </table> <i>bpgrp-id</i> — Specifies the bundle protection group identifier. <table> <tr> <td><i>bpgrp-type-bpgrp-num</i></td><td></td></tr> <tr> <td><i>bgrp</i></td><td>keyword</td></tr> <tr> <td><i>type</i></td><td>ima, ppp</td></tr> <tr> <td><i>bgrp-num</i></td><td>1 to 2000 (256 for 7750 SR-c12 and 7750 SR-c4)</td></tr> </table> <i>aps-id</i> — Specifies the APS group identifier. <table> <tr> <td><i>aps-group-id</i> [.channel]</td><td></td></tr> <tr> <td><i>aps</i></td><td>keyword</td></tr> <tr> <td><i>group-id</i></td><td>1 to 64 (16 for 7750 SR-c12 and 7750 SR-c4)</td></tr> </table> statistics — Specifies that port statistics will be cleared. atm — Specifies that ATM port statistics will be cleared. cp — Clears Connection Profile statistics. Values 1 to 8000 interface-connection — Clears interface-connection statistics. ilmi — Clears ILMI information. This parameter is only applicable for ports/bundles that support ILMI. pvc — Clears PVC statistics. <i>vpi</i> — Specifies the ATM network virtual path identifier (VPI) for this PVC. Values 0 to 4095 (NNI)	<i>bundle-type-slot/mda.bundle-num</i>		<i>bundle</i>	keyword	<i>type</i>	ima, fr, ppp	<i>bundle-num</i>	1 to 336	<i>bpgrp-type-bpgrp-num</i>		<i>bgrp</i>	keyword	<i>type</i>	ima, ppp	<i>bgrp-num</i>	1 to 2000 (256 for 7750 SR-c12 and 7750 SR-c4)	<i>aps-group-id</i> [.channel]		<i>aps</i>	keyword	<i>group-id</i>	1 to 64 (16 for 7750 SR-c12 and 7750 SR-c4)
<i>bundle-type-slot/mda.bundle-num</i>																							
<i>bundle</i>	keyword																						
<i>type</i>	ima, fr, ppp																						
<i>bundle-num</i>	1 to 336																						
<i>bpgrp-type-bpgrp-num</i>																							
<i>bgrp</i>	keyword																						
<i>type</i>	ima, ppp																						
<i>bgrp-num</i>	1 to 2000 (256 for 7750 SR-c12 and 7750 SR-c4)																						
<i>aps-group-id</i> [.channel]																							
<i>aps</i>	keyword																						
<i>group-id</i>	1 to 64 (16 for 7750 SR-c12 and 7750 SR-c4)																						

0 to 255 (UNI)

vci — Specifies the ATM network virtual channel identifier (VCI) for this PVC.

Values [1, 2, 5 to 65535] 1

pvp — Clears PVP statistics.

vpi — Specifies the ATM network virtual path identifier (VPI) for this PVP.

Values 0 to 4095 (NNI)
0 to 255 (UNI)

pvt — Clears PVT statistics.

vpi — Specifies the ATM network virtual path identifier (VPI) for this PVT.

Values [0 to 4095].[0 to 4095] (NNI)
[0 to 255].[0 to 255] (UNI)

slot — Specifies the slot number.

Values 1 to 10

mda — The MDA number.

Values 1, 2
MDA: 7750 SR-c12: 1, 3, 5, 7, 9, 11
CMA: 7750 SR-c12: 1 to 12

Default All MDAs

port-connection — Clears port-connection statistics.

phys-state-change-count — Clears the counter that tracks physical port state transitions for ethernet ports ("Phys State Chng Cnt" in "show port" output, or `tmnxPortPhysStateChangeCount` in the TIMETRA-PORT-MIB).

queue-group-name — Clears the specified port queue group name. It uniquely identifies a port ingress queue group in the managed system up to 32 characters.

instance-id — Specifies the specific instance of a queue group

Values 1 to 65535

ingress — Clears ingress queue group information.

egress — Clears egress queue group information.

access — Clears access mode statistics for the queue group.

network — Clears network mode statistics for the queue group.

ethernet — Specifies an Ethernet port will have the clear functions executed.

dampening — Clears the current accumulated penalties of the exponential port dampening feature for the port. The accumulated penalties can be cleared only with this command keyword.

efm-oam — Clears efm-oam statistics on the port.

events — Specifies an efm-oam event will be cleared.

local — Specifies that only local EFM OAM events will be cleared.

remote — Specifies that only remote (received from peer) events will be cleared. Local and remote is not specified.

Default Without specifying an option, both local and remote are cleared.

shaper-name — Clears information about the specified shaper name, up to 32 characters.

name — Clears information about the specified Vport name up to 32 characters.

standard value — Clears high water marks within specific standard port class pools using the standard parameters.

Values 1 to 6

alternate value — Clears high water marks within specific alternative port class pools using the alternative parameters.

Values 1 to 6

Output The following output is an example of port information.

Sample Output

```
A:SR12# clear port 3/1/1 atm
- atm

      cp          - Clear Connection Profile statistics
      ilmi        - Clear ILMI statistics
      interface-conn* - Clear interface-connection statistics
      pvc         - Clear PVC statistics
      pvp         - Clear PVP statistics
      pvt         - Clear PVT statistics

A:SR12# clear port 3/1/1 atm cp
- cp [<cp>] statistics

<cp>          : [1..8000]
<statistics>  :
```

2.21.2.11 Tools Commands

aps

Syntax **aps** *aps-id* [clear]
 aps mc-aps-signaling [clear]
 aps mc-aps-ppp [clear]

Context	tools>dump>aps
Description	This command displays Automated Protection Switching (APS) information.
Parameters	<i>aps-id</i> — Specifies the APS ID. Values <i>aps-group-id</i> - aps: keyword - group-id: 1 to 64 clear — Removes all Automated Protection Switching (APS) operational commands. mc-aps-signaling — Displays multi-chassis APS signaling information. mc-aps-ppp — Displays multi-chassis APS PPP information.
Output	The following output is an example of APS information.

Sample Output

```
*A:AS_SR7_2# tools dump aps aps-33

GrpId = 33, state = Running, mode:cfg/oper = Bi-directional/Bi-directional
  revert = 0, workPort: N/A, protPort: 2/1/1, activePort: working
  rxK1 = 0x0 (No-Req on Protect), physRxK1 = 0x0, rxK2 = 0x5
  txK1 = 0x0 (No-Req on Protect), physTxK1 = 0x0, txK2 = 0x5
  K1ReqToBeTxed = 0x0, K1ChanToBeTxed = 0x0, lastRxReq = 0xc
  MC-APS Nbr = 10.100.100.1 (Up), advIntvl = 10, hold = 30
  workPort: status = OK, Tx-Lais = None, sdCnt = 1, sfCnt = 1
    numSwitched = 1, switchSecs = 0, lastSwitched = 07/25/2007 08:00:12
    disCntTime = , alarms = , switchCmd = No Cmd
  protPort: status = OK, Tx-Lais = None, sdCnt = 1, sfCnt = 0
    numSwitched = 1, switchSecs = 0, lastSwitched = 07/25/2007 08:03:39
    disCntTime = , alarms = , switchCmd = No Cmd
  GrpStatus: OK, mmCnt = 1, cmCnt = 1, psbfCnt = 1, feplfCnt = 2
  LocalSwitchCmd: priority = No-Req, portNum = 0
  RemoteSwitchCmd: priority = No-Req, portNum = 0
  Running Timers = mcAdvIntvl mcHold
  processFlag = apsFailures = , sonet = Y
  DebugInfo: dmEv = 0, dmClrEv = 0, amEv = 1, amClrEv = 1
    cmEv = 1, cmClrEv = 1, psbfEv = 1, psbfClrEv = 1
    feplfEv = 2, feplfClrEv = 2, wtrEv = 0, psbfDetectEv = 0
    wSdEv = 1, wSfEv = 2, pSdEv = 1, pSfEv = 1
    portStatusEv = 8, rxK1Ev = 9, txLaisEv = 2, lastEvName = FeplClr
    CtlUpEv = 3, CtlDnEv = 2, wAct = 0, wDeAct = 0

Seq      Event  TxK1/K2 RxK1/K2  Dir    Active    Time
===      =====
000      ProtAdd 0xc005  0x0000  Tx-->   Work    497 02:18:10.590
001      RxKByte 0xc005  0x6dea  Rx<--   Work    497 02:20:14.820
002      RxKByte 0xc005  0xc005  Rx<--   Work    497 02:21:30.970
003      RxKByte 0xc005  0x2005  Rx<--   Work    497 02:21:36.530
004      pSFClr  0x0005  0x2005  Tx-->   Work    497 02:21:40.590
005      RxKByte 0x0005  0x0005  Rx<--   Work    497 02:21:40.600
006      RxKByte 0x0005  0xc115  Rx<--   Work    497 02:25:22.840
007      RxKByte 0x2115  0xc115  Tx-->   Prot    497 02:25:22.840
008      RxKByte 0x2115  0xa115  Rx<--   Prot    000 00:00:47.070
009      RxKByte 0x2115  0x1115  Rx<--   Prot    000 00:00:47.560
010      RxKByte 0x2115  0xc005  Rx<--   Prot    000 00:00:57.010
```

```
011      RxKByte 0x2005  0xc005  Tx-->      Work  000 00:00:57.010
012      RxKByte 0x2005  0x0005  Rx<--      Work  000 00:01:06.170
013      RxKByte 0x0005  0x0005  Tx-->      Work  000 00:01:06.170
```

Sample Output

```
:AS_SR7_1# tools dump aps mc-aps-ppp

pppmMcsModStarted = Yes
pppmMcsDbgDoSync = Yes
pppmMcsApsGrpHaAuditDone = Yes
pppmMcsPostHaSyncedApsGrpId = 47
pppmMcsMcApsChanCnt = 1280

pppmMcsDbgRxPktCnt = 2560
pppmMcsDbgRxPktNotProcessedCnt = 0
pppmMcsDbgRxPktInvalidCnt = 0
pppmMcsDbgInconsistentRxPktDropCnt = 0
pppmMcsDbgInconsistentTxPktDropCnt = 1176
pppmMcsDbgTxPktNotSentCnt = 0
pppmMcsDbgTxPktSentCnt = 25
pppmMcsDbgEvtDropCnt = 0
pppmMcsDbgMemAllocErrCnt = 0
pppmMcsDbgReTxCnt = 0
pppmMcsDbgReTxExpCnt = 0
pppmMcsDbgReReqCnt = 0

pppmMcsStateAckQueueCnt (curr/peek) = 0/130
pppmMcsStateReqQueueCnt (curr/peek) = 0/1280
pppmMcsStateReReqQueueCnt (curr/peek) = 0/256
pppmMcsStateTxQueueCnt (curr/peek) = 0/512
pppmMcsStateReTxQueueCnt (curr/peek) = 0/130

MC-APS Peer Info :
-----

Grp 13 Addr 10.100.100.2 - Up
Grp 20 Addr 10.100.100.2 - Up
Grp 35 Addr 10.100.100.2 - Up
Grp 43 Addr 10.100.100.2 - Up
Grp 47 Addr 10.100.100.2 - Up

Number of ppmMcs Evt Msgs dispatched:
  ctl_link_state : 0
  ctl_link_up_tmr : 0
  ctl_link_down_tmr : 0
  ha_audit_done : 0
```

Sample Output

```
*A:eth_aps_sr7# tools dump aps mc-aps-signaling

MC-APS Control Debug Counters :
-----

Ctl Pkt Rx = 0
Invalid Rx Ctl Pkt = 0
Incompatible Rx Ctl Pkt = 0
```

```
Nbr not Rx Ctl Pkt = 0
Invalid Rx Ctl Pkt Tlv = 0
Ctl Pkt Rx-ed before HaReady = 0
Not sent Tx Ctl Pkt = 0

MC-APS-LAG Debug Counters :
-----
Ctl Pkt Rx from IOM          = 0

Not processed Rx Ctl Pkt     = 0
Invalid Rx Ctl Pkt           = 0
Incompatible Rx Ctl Pkt      = 0
Rx Ctl Pkt queueing failed = 0

Ctl Pkt Tx (direct)          = 0
Ctl Pkt Tx (UDP socket)      = 0
Not sent Tx Ctl Pkt          = 0

Route Update                  = 0
Matched Route Update          = 0

Msg Buf Alloc Failed          = 0

MC-APS-LAG NbrRoute Entries :
-----
NbrAddr 10.1.1.1 NextHopAddr ::
  EgressIfIndex = 0
  EgressPortId  = Unknown
  app refCnt     = 1
  refCntTotal    = 1
```

aps

Syntax	aps
Context	tools>perform
Description	This command enables the context to perform Automated Protection Switching (APS) operations.

clear

Syntax	clear <i>aps-id</i> { protect working } [number <i>number</i>]				
Context	tools>perform>aps				
Description	This command removes all Automated Protection Switching (APS) operational commands.				
Parameters	<i>aps-id</i> — This option clears a specific APS on un-bundled SONET/SDH ports.				
Values	<table><tr><td><i>aps-id</i></td><td><i>aps-group-id</i></td></tr><tr><td>aps</td><td>keyword</td></tr></table>	<i>aps-id</i>	<i>aps-group-id</i>	aps	keyword
<i>aps-id</i>	<i>aps-group-id</i>				
aps	keyword				

group-id 1 to 128

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

number — Specifies the APS channel number.

Values 1, 2

clear-lockout-annexb

Syntax clear-lockout-annexb *aps-id*

Context tools>perform>aps

Description This command clears lockout of the Annex B APS group.

Parameters *aps-id* — Specifies an APS ID.

Values

<i>aps-id</i>	<i>aps-group-id</i>
aps	keyword
group-id	1 to 128

exercise

Syntax exercise *aps-id* {**protect** | **working**}

Context tools>perform>aps

Description This command performs an exercise request on the protection or working circuit.

Parameters *aps-id* — Specifies the APS ID.

Values

<i>aps-id</i>	<i>aps-group-id</i>
aps	keyword
group-id	1 to 128

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

force

Syntax	force <i>aps-id</i> { protect working } [number <i>number</i>]
Context	tools>perform>aps
Description	This command forces a switch to either the protect or working circuit
Parameters	<i>aps-id</i> — Specifies the APS ID.

Values

<i>aps-id</i>	<i>aps-group-id</i>
aps	keyword
group-id	1 to 128

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

number — Specifies the APS channel number.

Values 1, 2

lockout

Syntax	lockout <i>aps-id</i>
Context	tools>perform>aps
Description	This command locks out the protection circuit.
Parameters	<i>aps-id</i> — Specifies the APS ID.

Values

<i>aps-id</i>	<i>aps-group-id</i>
aps	keyword
group-id	1 to 128

lockout-annexb

Syntax	lockout-annexb <i>aps-id</i>
Context	tools>perform>aps
Description	This command locks out the Annex B APS group.

Parameters	<i>aps-id</i> — Specifies the APS ID.		
Values	<i>aps-id</i>	<i>aps-group-id</i>	
		aps	keyword
		group-id	1 to 128

request

Syntax	request <i>aps-id</i> {protect working}						
Context	tools>perform>aps						
Description	This command requests a manual switch to protection or working circuit.						
Parameters	<i>aps-id</i> — Specifies the APS ID.						
Values	<table><tr><td><i>aps-id</i></td><td><i>aps-group-id</i></td></tr><tr><td>aps</td><td>keyword</td></tr><tr><td>group-id</td><td>1 to 128</td></tr></table>	<i>aps-id</i>	<i>aps-group-id</i>	aps	keyword	group-id	1 to 128
<i>aps-id</i>	<i>aps-group-id</i>						
aps	keyword						
group-id	1 to 128						

protect — This command requests a manual switch to a port that is acting as the protection circuit for the APS group.

working — This command requests a manual switch to a port that is acting as the working circuit for this APS group.

eth-ring

Syntax	eth-ring
Context	tools>perform
Description	This command performs Ethernet ring operations.

clear

Syntax	clear <i>ring-index</i>
Context	tools>perform>eth-ring
Description	This command, at the Ethernet Ring Node, is used for the following operations: a. Clearing an active local administrative command (for example, Forced Switch or Manual Switch).

- b. Triggering reversion before the WTR or WTB timer expires in case of revertive operation.
- c. Triggering reversion in case of non-revertive operation.

Parameters *ring-index* — This option clears a specific Ethernet ring.

Values 1 to 128

force

Syntax **force** *ring-index* **path** {a | b}

Context tools>perform>eth-ring

Description This command forces a block on the ring port where the command is issued.

Parameters *ring-index* — Specifies the ring index.

Values 1 to 128

path — Displays information for a specific path.

Values a, b

manual

Syntax **manual** *ring-index* **path** {a | b}

Context tools>perform>eth-ring

Description In the absence of a failure or FS, this command forces a block on the ring port where the command is issued.

Parameters *ring-index* — Specifies the ring index.

Values 1 to 128

path — Displays information for a specific path.

Values a, b

eth-tunnel

Syntax **eth-tunnel** *tunnel-index* [**clear**]

Context tools>dump

Description This command displays Ethernet tunnel information.

Parameters *tunnel-id* — Specifies the tunnel ID.

Values 1 to 128

clear — Clears statistics after reading.

Output The following output is an example of Ethernet tunnel information.

Sample Output

```
*A:PE-E# tools dump eth-tunnel 1

TunnelId 1 (Up/Up), Port eth-tunnel-1 (Up/Up): type g8031-1to1
NumMems 2/2, Up/Dn 0/0, active 0x1, present 0x3 baseMemPort 1/1/2
  memId 1 (P), port 1/1/2 (Up), tag 1.0(Up) status (Up/Up)
    ccCnt-sf/ok 1/1 idx 0 tunId 1
  memId 2 (S), port 2/1/2 (Up), tag 1.0(Up) status (Up/Up)
    ccCnt-sf/ok 0/0 idx 1 tunId 1

TunId = 1, state = Running, Active = Work, Now = 000 00:16:48.140
revert = 1, ReqState = NR-NULL, Pdu(Tx/Rx): 0x0f0000/0x0f0000
Defects =
Running Timers = PduReTx
  Work MemId = 1 (1/1/2:1.0), state = Ok, cc = 000 00:16:23.510U
    ActiveCnt = 4, ActiveSeconds = 791
  Protect MemId = 2 (2/1/2:1.0), state = Ok, cc = 000 00:09:47.560U
    ActiveCnt = 3, ActiveSeconds = 308
DbgCnts: swEv = 2, wMemSts = 2, pMemSts = 0
  rxPdu (valid/invalid) = 4/0, wSfClr = 1, pSfClr = 0, wtrExp = 1
  cm = 0, cmClr = 0, pm = 0, pmClr = 0, nr = 0, nrClr = 0
Seq  Event      TxPdu      RxPdu      Dir      Act      Time
===  =====
000  wMemSts 0xbf0101 wSF 0x0f0000 NR Tx--> Prot 000 00:16:12.450
001      RxPdu 0xbf0101 wSF 0x0f0101 NR Rx<-- Prot 000 00:16:12.450
002      RxPdu 0xbf0101 wSF 0xbf0101 wSF Rx<-- Prot 000 00:16:12.480
003      RxPdu 0xbf0101 wSF 0x0f0101 NR Rx<-- Prot 000 00:16:24.890
004  wSFClr 0x5f0101 WTR 0x0f0101 NR Tx--> Prot 000 00:16:25.030
005      WTR 0x0f0000 NR 0x0f0101 NR Tx--> Work 000 00:16:26.630
006      RxPdu 0x0f0000 NR 0x0f0000 NR Rx<-- Work 000 00:16:26.630
*A:PE-E#
```

frame-relay

Syntax **frame-relay** *port-id*

Context tools>dump

Description This command displays frame-relay information.

Parameters *port-id* — Specifies the physical port ID.

port-id *slot/mda/port* [*.channel*]
 bundle-id *bundle-type-slot/mda.bundle-num*
 bundle *keyword*

	type	fr
	bundle-num	1 to 336
eth-sat-id	<i>esat-id/slot/port</i>	
	esat	keyword
	<i>id</i>	1 to 20
pxc-id	<i>pxc-id.sub-port</i>	
	pxc	keyword
	<i>id</i>	1 to 64
	<i>sub-port</i>	a, b

Output The following output is an example of frame relay information.

Sample Output

```
tools dump frame-relay bundle-fr-1/1.1
```

```
MLFR info for bundle-fr-1/1.1:
  expectedBundleIdStr : "7750-bundle-fr-1/1.1"
  bundleId            : 1
  active/configured   : 4/4
  adminEnabled        : 1
  ddDebDownCnt        : 10
  ddMaxForBundle      : 0
  fwdId               : 0
  linkDebugMask       : 0

----- Member Information -----

1/1/1.1.1.1          PortId = 0x22208056
  Internal ID         : bundle 1, link 1
  Link protocol state : Up
  Diff delay state    : Yellow
  ddRedCnt            : 0
  ddYellowCnt         : 11
  Smoothed diff delay : 20 ms.
  Historical RTT       : 22788, 22756, 22752 us.
  Rx LinkId           : "1/1/1.1.1.1"
  WARNING             : dbgDelay=20 ms
  LIP CTL Messages    : Tx: 136      Rx: 135
  Hello Messages      : Tx: 121      Rx: 12
  Hello Ack Messages  : Tx: 12       Rx: 121
  AddLink Messages    : Tx: 2        Rx: 1
  AddLinkAck Messages : Tx: 1        Rx: 1
  AddLinkRej Messages : Tx: 0        Rx: 0

1/1/1.1.2.1          PortId = 0x22208077
  Internal ID         : bundle 1, link 2
  Link protocol state : Up
  Diff delay state    : OK
  Smoothed diff delay : 0 ms.
  Historical RTT       : 2271, 2304, 2309 us.
  Rx LinkId           : "1/1/1.1.2.1"
```

```

LIP CTL Messages      : Tx: 136      Rx: 135
Hello Messages        : Tx: 121      Rx: 12
Hello Ack Messages    : Tx: 12       Rx: 121
AddLink Messages      : Tx: 2        Rx: 1
AddLinkAck Messages   : Tx: 1        Rx: 1
AddLinkRej Messages   : Tx: 0        Rx: 0
1/1/1.1.3.1           PortId = 0x22208098
Internal ID           : bundle 1, link 4
Link protocol state    : Up
Diff delay state       : OK
Smoothed diff delay    : 0 ms.
Historical RTT         : 2218, 2223, 2359 us.
Rx LinkId              : "1/1/1.1.3.1"
LIP CTL Messages      : Tx: 14       Rx: 13
Hello Messages        : Tx: 8        Rx: 1
Hello Ack Messages    : Tx: 1        Rx: 8
AddLink Messages      : Tx: 3        Rx: 1
AddLinkAck Messages   : Tx: 1        Rx: 2
AddLinkRej Messages   : Tx: 0        Rx: 0
1/1/1.1.4.1           PortId = 0x222080b9
Internal ID           : bundle 1, link 3
Link protocol state    : Up
Diff delay state       : OK
Smoothed diff delay    : 0 ms.
Historical RTT         : 2248, 2242, 2309 us.
Rx LinkId              : "1/1/1.1.4.1"
LIP CTL Messages      : Tx: 14       Rx: 13
Hello Messages        : Tx: 8        Rx: 1
Hello Ack Messages    : Tx: 1        Rx: 8
AddLink Messages      : Tx: 3        Rx: 1
AddLinkAck Messages   : Tx: 1        Rx: 2
AddLinkRej Messages   : Tx: 0        Rx: 0

```

lag

Syntax **lag lag-id lag-id**

Context tools>dump

Description This command dumps LAG information.

Parameters *lag-id* — Specifies the LAG ID.

Values 1 to 800

map-to-phy-port

Syntax **map-to-phy-port {ccag ccag-id | lag lag-id | eth-tunnel tunnel-index} {isid isid [end-isid isid] | service service-id | svc-name [end-service service-id | svc-name]} [summary]**

Context tools>dump

Description	This command provides the ability to respond to a query to provide the link in a LAG/Ethernet tunnel (load sharing protection mode)/CCAG that is currently assigned to a given service-id or ISID.
Parameters	<i>lag-id</i> — Specifies the LAG ID. Values 1 to 800 <i>isid</i> — Specifies the ISID. Values 0 to 16777215 <i>service-id</i> — Specifies the service ID. Values 1 to 2148278381, svc-name: 64 char max <i>tunnel-index</i> — Specifies the tunnel index. Values 1 to 1024 <i>ccag-id</i> — Specifies the CCAG ID. Values 1 to 8 summary — Displays summary information.

port

Syntax	port <i>port-id</i>		
Context	tools>dump		
Description	This command enables the context to dump port information.		
Parameters	<i>port-id</i> — Specifies the physical port ID.		
	<i>port-id</i>	<i>slot/mda/port</i>	
	bundle-id	bundle-type-slot/mda.bundle-num	
		bundle	keyword
		<i>type</i>	ima, fr, ppp
		<i>bundle-num</i>	1 to 336
	bpgrp-id	bpgrp-type-bpgrp-num	
		bpgrp	keyword
		<i>type</i>	ima, ppp
		<i>bpgrp-num</i>	1 to 2000
	aps-id	aps-group-id	
		aps	keyword
		<i>group-id</i>	1 to 128

dwdm

Syntax	dwdm
Context	tools>dump>port
Description	This command enables the context to dump Dense Wavelength Division Multiplexing (DWDM) information for the port.

coherent

Syntax	coherent
Context	tools>dump>port>dwdm
Description	This command enables the context to dump coherent optics information for the port.

cpr-wndw-sz-srch-info

Syntax	cpr-wndw-sz-srch-info
Context	tools>dump>port>dwdm>coherent
Description	This command dumps the Carrier Phase Recovery window size search information for the port.

pcs

Syntax	pcs[clear]
Context	tools>dump>port
Description	This command dumps the Physical Coding Sublayer (PCS) information for the port.

rs-fec

Syntax	rs-fec [clear]
Context	tools>dump>port
Description	This command dumps the Reed-Solomon Forward Error Correction (RS-FEC) information for port.

ppp

Syntax	ppp <i>port-id</i>		
Context	tools>dump		
Description	This command displays PPP information for a port.		
Parameters	<i>port-id</i> — Specifies the physical port ID.		
	<i>port-id</i>	<i>slot/mda/port</i> [<i>channel</i>]	
	bundle-id	bundle-type-slot/mda.bundle-num	
		bundle	keyword
		type	ppp
		bundle-num	1 to 336
	bpgrp	bpgrp-type-bpgrp-num	
		bpgrp	keyword
		type	ppp
		bpgrp-num	1 to 2000
	aps-id	aps-group-id[channel.]	
		aps	keyword
		group-id	1 to 128
	eth-sat-id	esat-id/slot/port	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-id.sub-port	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

Output The following output is an example of PPP information.

Sample Output

```
*A:sr7# tools dump ppp aps-1.1.1.1
=====
Id           : aps-1.1.1.1      ppp unit       : 40
member of    : bpgrp-ppp-1
=====
looped back  : no              dbgMask        : 0x0
-----
LCP
-----
phase        : NETWORK         state         : OPENED
passive      : off             silent        : off
restart      : on
```

```

mru          : 1500          mtu          : 1502
ack'd peer mru : 1500
got local mrru : 1524
local magic   : 0x0          peer magic   : 0x0

keepalive    : on           echo num     : 2
echo timer   : on           echos fail   : 3
echo intv    : 10           echos pend   : 0

options      mru      asyncMap upap      chap      magic      pfc
we negotiate Yes      No      No      No      No      Yes
peer ack'd   Yes      No      No      No      No      No
we allow     Yes      No      No      No      No      Yes
we ack'd     Yes      No      No      No      No      No

options      acfc      lqr      mrru      shortSeq endPoint mlhdrfmt
we negotiate Yes      No      Yes      No      Yes      No
peer ack'd   No      No      Yes      No      Yes      No
we allow     Yes      No      Yes      Yes      Yes      No
we ack'd     No      No      Yes      No      Yes      No
=====
*A: sr7#

```

redundancy

Syntax	redundancy
Context	tools>dump
Description	This command enables the context to dump redundancy parameters.

multi-chassis

Syntax	multi-chassis
Context	tools>dump>redundancy
Description	This command enables the context to dump multi-chassis parameters.

mc-ring

Syntax	mc-ring mc-ring peer <i>ip-address</i> [ring sync-tag]
Context	tools>dump>redundancy>multi-chassis
Description	This command dumps multi-chassis ring data.
Parameters	<i>ip-address</i> — Specifies the peer's IP address.

sync-tag — Specifies the ring's sync-tag created in the **config>redundancy>mc>peer>mcr>ring** context.

srrp-sync-database

Syntax	srrp-sync-database [instance <i>instance-id</i>] [peer <i>ip-address</i>]
Context	tools>dump>redundancy>multi-chassis
Description	This command dumps multi-chassis SRRP sync database information.
Parameters	<i>instance-id</i> — Specifies the instance ID. Values 1 to 4294967295 <i>ip-address</i> — Dumps the specified address (in the form of a.b.c.d). Values ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d, where x: [0 to FFFF]H and d: [0 to 255]D

sync-database

Syntax	sync-database [peer <i>ip-address</i>] [port <i>port-id</i> <i>lag-id</i>] [sync-tag <i>sync-tag</i>] [application <i>application</i>] [detail] [type <i>type</i>] sync-database [peer <i>ip-address</i>] [sdp <i>sdp-id</i>] [sync-tag <i>sync-tag</i>] [application <i>application</i>] [detail] [type <i>type</i>]
Context	tools>dump>redundancy>multi-chassis
Description	This command dumps multi-chassis sync database information.
Parameters	<i>ip-address</i> — Dumps the specified address of the multi-chassis peer. Values ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d • x: [0 to FFFF]H • d: [0 to 255]D <i>port-id</i> — Dumps the specified port ID of the multi-chassis peer. Values <i>port-id, lag-id</i> slot/mda/port <i>lag-id</i> lag-id

	lag	keyword
	id	1 to 800
pw-id	pw-id	
	pw	keyword
	id	1 to 32767

lag-id — Dumps the specified Link Aggregation Group (LAG) on this system.

Values

<i>port-id, lag-id</i>	slot/mda/port	
	lag-id	lag-id
	lag	keyword
	id	1 to 800
pw-id	pw-id	
	pw	keyword
	id	1 to 32767

sync-tag — Dumps the synchronization tag used while synchronizing this port with the multi-chassis peer up to 32 characters in length.

application — Dumps the specified application information that was synchronized with the multi-chassis peer.

Values

dhcp-server	local DHCP server
igmp	internet group management protocol
igmp-snooping	IGMP snooping
mc-ring	multi-chassis ring
l2tp	L2TP
mld	multicast listener discovery
mld-snooping	multicast listener discovery snooping
srrp	simple router redundancy protocol
sub-host-trk	subscriber host tracking
sub-mgmt-ipoe	subscriber management for IPoE
sub-mgmt-pppoe	subscriber management for PPPoE
mc-ipsec	multi-chassis IPsec
python	Python cache
diameter-proxy	diameter proxy
pim-snpg-sap	protocol independent multicast snooping for SAP
pim-snpg-sdp	protocol independent multicast snooping for SDP

detail — Displays detailed information.

type — Displays information for the specified type.

Values alarm-deleted, local-deleted, global-deleted, omcr-standby, omcr-alarmed

sdp-id — Displays information for the specified SDP ID.

Values 1 to 32767

Output The following output is an example of sync database information.

Sample Output

```
A:Dut-C# tools dump redundancy multi-chassis sync-database application
```

```
<ip-address>      : a.b.c.d
<port-id|lag-id>  : slot/mda/port or lag-<lag-id>
<sync-tag>        : [32 chars max]
<application>     : dhcp-server      - local dhcp server
                   igmp              - internet group management protocol
                   igmp-snooping     - igmp-snooping
                   mc-ring           - multi-chassis ring
                   mld               - multicast listener discovery
                   mld-snooping      - multicast listener discovery-snooping
                   srrp              - simple router redundancy protocol
                   sub-host-trk      - subscriber host tracking
                   sub-mgmt-ipoe     - subscriber management for IPoE
                   sub-mgmt-pppoe    - subscriber management for PPPoE
                   mc-ipsec          - multi-chassis IPsec
<detail>          : - displays detailed information
<type>            : alarm-deleted|local-deleted|global-deleted|
                   omcr-standby|omcr-alarmed
```

ima

Syntax ima

Context tools>perform

Description This command allows the use of IMA operations.

reset

Syntax reset *bundle-id*

Context tools>perform>ima

Description This command sets an IMA-bundle to the Start Up state.

Parameters	<i>bundle-id</i> — Specifies an existing bundle ID.
Values	bundle-ima-slot/mda. <i>bundle-num</i> bundle-num: 1 to 256

lag

Syntax	lag
Context	tools>perform
Description	This command provides tools for controlling LAG.

clear-force

Syntax	clear-force all-mc clear-force lag-id <i>lag-id</i> [sub-group <i>sub-group-id</i>] clear-force peer-mc <i>ip-address</i>
Context	tools>perform>lag
Description	This command clears forced status.
Parameters	all-mc — Specifies all MC-LAGs. <i>lag-id</i> — Specifies the LAG ID. Values 1 to 800 <i>sub-group-id</i> — Specifies the subscriber group ID. Values 1 to 16 <i>ip-address</i> — Specifies the peer MC IP address. Values ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d, where x: [0 to FFFF]H and d: [0 to 255]D

force

Syntax	force all-mc {active standby} force lag-id <i>lag-id</i> [sub-group <i>sub-group-id</i>] {active standby} force peer-mc <i>ip-address</i> {active standby}
Context	tools>perform>lag

Description	This command allows forcing the specified LAG, subgroup, all MC-LAGs, or remote peer for MC-LAGs to become active or standby when LAG runs in Active/Standby mode. To remove the forced condition, execute the tools perform lag clear-force command.
Parameters	all-mc — Specifies all MC-LAGs. active — Specifies to become active. standby — Specifies to become standby. <i>lag-id</i> — Specifies the LAG ID. Values 1 to 800 <i>sub-group-id</i> — Specifies the subscriber group ID. Values 1 to 16 <i>ip-address</i> — Specifies the IP address. Values ipv4-address: a.b.c.d ipv6-address: • x:x:x:x:x:x:x:x (eight 16-bit pieces) • x:x:x:x:x:d.d.d.d, where x: [0 to FFFF]H and d: [0 to 255]D

load-balance

Syntax	load-balance lag-id lag-id [class {1 2 3}]
Context	tools>perform>lag
Description	Load balance specified LAG's links when per-link-hash weighted is deployed. Load balancing can be per specified class or on all classes if no class is specified.
Parameters	<i>lag-id</i> — Specifies the LAG ID. Values 1 to 800 class — Specifies the class. Values 1, 2, 3

2.21.2.11.1 Debug Commands

lmi

Syntax	lmi [port-id] no lmi
Context	debug>frame-relay

Description This debug command enables tracing of all the LMI messages in both receive and transmit directions for one or all of the Frame Relay interfaces. All types of Frame Relay interfaces are supported. If the port ID is not specified, debug is enabled on all Frame Relay interfaces.

The **no** form of the command turns off LMI and Frame-Relay debugging, **debug>frame-relay>no lmi** and **debug>no frame-relay**.

Parameters *port-id* — Specifies the ILMI-supporting port ID.

Values

<i>port-id</i>	<i>slot/mda/port [.channel]</i>		
	eth-sat-id	esat-id/slot/port	
		esat	keyword
		id	1 to 20
	pxc-id	pxc-id.sub-port	
		pxc	keyword
		id	1 to 64
		sub-port	a, b

Output The following output is an example of frame relay LMI information.

Sample Output

```
2959 2007/04/11 23:01:34.63 MINOR: DEBUG #2001 - FR
"FR: TX STATUS Msg on dce Port: 1/1/1 LMI: itu
FR Hdr: 00 01 03 08 00 7D
Rpt IE: 51 01 01 LINK_INT_VERIFY
KA IE: 53 02 31 45 TxSeqNo=49 RxSeqNo=69"

2960 2007/04/11 23:01:44.63 MINOR: DEBUG #2001 - FR
"FR: RX STATUS ENQ Msg on dce Port: 1/1/1 LMI: itu
FR Hdr: 00 01 03 08 00 75
Rpt IE: 51 01 01 LINK_INT_VERIFY
KA IE: 53 02 46 31 TxSeqNo=70 RxSeqNo=49"
```

atm

Syntax [no] atm

Context debug

Description This command enables, disables and configures debugging for the ATM.

ilmi

Syntax ilmi [*port-id*]

	no ilmi <i>port-id</i>
Context	debug>atm
Description	This command enables debugging for ATM ILMI. The no form of the command turns off ILMI and debugging.
Parameters	<i>port-id</i> — Specifies the port ID. Values slot/mda/port [.channel] or bundle-ima-slot/mda. bundle-num bundle: keyword ima: keyword bundle-num: 1 to 56
Output	The following output is an example of ATM ILMI information

Sample Output

```

A:CHRISILMI# debug atm no ilmi 1/2/2
A:CHRISILMI# debug atm ilmi 1/2/4

In kernel:
ILMI_DEBUG_LOG {557907970}: 21:32:28
PDU DUMP (RAW):
 30 77 02 01 00 04 04 49 4c 4d 49 a4
 6c 06 07 2b 06 01 04 01 82 61 40 04
 00 00 00 00 02 01 06 02 01 02 43 03
 1b 24 b1 30 50 30 12 06 0d 2b 06 01
 04 01 82 61 02 05 01 01 01 00 02 01
 00 30 12 06 0d 2b 06 01 04 01 82 61
 02 05 01 01 03 00 02 01 21 30 12 06
 0d 2b 06 01 04 01 82 61 02 05 01 01
 02 00 02 01 00 30 12 06 0d 2b 06 01
 04 01 82 61 02 05 01 01 04 00 02 01

00 00 00
PDU DUMP (DECODED):
PDU Length: 123
Community: ILMI
Version: 1
Msg Type: SNMP TRAP
ObjectId: 1.3.6.1.4.1.353.2.5.1.1.1.0
ObjectId: 1.3.6.1.4.1.353.2.5.1.1.3.0
ObjectId: 1.3.6.1.4.1.353.2.5.1.1.2.0
ObjectId: 1.3.6.1.4.1.353.2.5.1.1.4.0
...
```

cisco-hdlc

Syntax **cisco-hdlc** [*port-id*]
no cisco-hdlc

Context	debug
Description	This command configures debugging for Cisco-HDLC encapsulation.
Parameters	<i>port-id</i> — Specifies the physical port ID.
Values	slot/mda/port[.channel]

frame-relay

Syntax	[no] frame-relay
Context	debug
Description	This command enables, disables and configures debugging for frame relay.

frf16

Syntax	[no] frf16 <i>port-id</i>
Context	debug>frame-relay
Description	This command enables tracing of all FRF16 compliant MLFR link integrity protocol messages in both the receive and transmit directions on a specific member link of an MLFR bundle. The no form of the command turns off MLFR debugging.
Parameters	<i>port-id</i> — Specifies the port ID of the FRF16 bundle member link. <i>port-id</i> <i>slot/mda/port [.channel]</i> <i>bundle-id</i> <i>bundle-type-slot/mda.bundle-num</i> <i>bundle</i> <i>keyword</i> <i>type</i> <i>fr</i> <i>bundle-num</i> 1 to 336

Output The following output is an example of frame relay FRF16 information.

Sample Output

```

1 2009/02/18 10:39:42.74 UTC MINOR: DEBUG #2001 Base MLFR
"MLFR: [_LIP_ParseRxFrame]
RxMsg <bundle-fr-1/1.1:1/1/1.0x56> <state-0: Up>
<MsgType-5:HelloAck>
    <ie-03:      MagicNum>,<len 06>,<27002>
    <ie-05:      TimeStamp>,<len 06>,<0x4b1c4558>"

2 2009/02/18 10:39:43.73 UTC MINOR: DEBUG #2001 Base MLFR
"MLFR: [_LIP_TxFrame]
TxMsg <bundle-fr-1/1.1:1/1/1.0x56> <state-0: Up>
<MsgType-4:Hello>

```

```

        <ie-03:      MagicNum>,<len 06>, <31104>
        <ie-05:      TimeStamp>,<len 06>, <0x5d804569>"

3 2009/02/18 10:39:43.73 UTC MINOR: DEBUG #2001 Base MLFR
"MLFR: [_LIP_ParseRxFrame]
RxMsg <bundle-fr-1/1.1:1/1/1.0x56> <state-0: Up>
<MsgType-5:HelloAck>
        <ie-03:      MagicNum>,<len 06>, <27002>
        <ie-05:      TimeStamp>,<len 06>, <0x5d804569>"

```

lmi

Syntax	lmi [<i>port-id</i>] no lmi
Context	debug>frame-relay
Description	This debug command enables tracing of all the LMI messages in both receive and transmit directions for one or all of the Frame Relay interfaces. All types of Frame Relay interfaces are supported. If the port ID is not specified, debug is enabled on all Frame Relay interfaces. The no form of the command turns off LMI and Frame-Relay debugging, debug>frame-relay>no lmi and debug>no frame-relay .
Parameters	<i>port-id</i> — Specifies the LMI-supporting port ID. Values <i>slot/mda/port[.channel]</i>
Output	The following output is an example of LMI information.

Sample Output

```

2959 2007/04/11 23:01:34.63 MINOR: DEBUG #2001 - FR
"FR: TX STATUS Msg on dce Port: 1/1/1 LMI: itu
FR Hdr: 00 01 03 08 00 7D
Rpt IE: 51 01 01 LINK_INT_VERIFY
KA IE: 53 02 31 45 TxSeqNo=49 RxSeqNo=69"

2960 2007/04/11 23:01:44.63 MINOR: DEBUG #2001 - FR
"FR: RX STATUS ENQ Msg on dce Port: 1/1/1 LMI: itu
FR Hdr: 00 01 03 08 00 75
Rpt IE: 51 01 01 LINK_INT_VERIFY
KA IE: 53 02 46 31 TxSeqNo=70 RxSeqNo=49"

```

lag

Syntax	lag [<i>lag-id lag-id</i> [<i>port port-id</i>]] [all] lag [<i>lag-id lag-id</i> [<i>port port-id</i>]] [sm] [pkt] [cfg] [red] [iom-upd] [port-state] [timers] [sel-logic] [mc] [mc-pkt] no lag [<i>lag-id lag-id</i>]

Context	debug
Description	This command enables debugging for LAG.
Parameters	<i>lag-id</i> — Specifies the link aggregation group ID. Values 1 to 800 <i>port-id</i> — Specifies the physical port ID. Values <i>slot/mda/port</i> all — Specifies to display all LAG information. sm — Specifies to display trace LACP state machine. pkt — Specifies to display trace LACP packets. cfg — Specifies to display trace LAG configuration. red — Specifies to display trace LAG high availability. iom-upd — Specifies to display trace LAG IOM updates. port-state — Specifies to display trace LAG port state transitions. timers — Specifies to display trace LAG timers. sel-logic — Specifies to display trace LACP selection logic. mc — Specifies to display multi-chassis parameters. mc-packet — Specifies to display the MC-LAG control packets with valid authentication were received on this system.

ppp

Syntax	ppp <i>port-id</i> no ppp [<i>port-id</i>]														
Context	debug														
Description	This command enables/disables and configures debugging for PPP.														
Parameters	<i>port-id</i> — Specifies the physical port ID. Values <table> <tr> <td><i>port-id</i></td><td><i>slot/mda/port</i> [<i>channel</i>]</td></tr> <tr> <td>bundle ID</td><td>bundle-type-slot/mda.<i>bundle-num</i></td></tr> <tr> <td></td><td>bundle keyword</td></tr> <tr> <td></td><td><i>type</i> ppp</td></tr> <tr> <td></td><td><i>bundle-num</i> 1 to 336</td></tr> <tr> <td>bpggrp-id</td><td>bpggrp-type-bpggrp-num</td></tr> <tr> <td></td><td>bpggrp keyword</td></tr> </table>	<i>port-id</i>	<i>slot/mda/port</i> [<i>channel</i>]	bundle ID	bundle-type-slot/mda. <i>bundle-num</i>		bundle keyword		<i>type</i> ppp		<i>bundle-num</i> 1 to 336	bpggrp-id	bpggrp-type-bpggrp-num		bpggrp keyword
<i>port-id</i>	<i>slot/mda/port</i> [<i>channel</i>]														
bundle ID	bundle-type-slot/mda. <i>bundle-num</i>														
	bundle keyword														
	<i>type</i> ppp														
	<i>bundle-num</i> 1 to 336														
bpggrp-id	bpggrp-type-bpggrp-num														
	bpggrp keyword														

aps-id	type	ppp
	bpgrp-num	1 to 2000
	aps-group-id[.channel]	
	aps	keyword
	group-id	1 to 128

3 Standards and Protocol Support



Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect (localised ID)*

draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*
RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*
RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/
MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual
Private Networks (VPNs)*
RFC 4724, *Graceful Restart Mechanism for BGP (helper mode)*
RFC 4760, *Multiprotocol Extensions for BGP-4*
RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge
Routers (6PE)*
RFC 4893, *BGP Support for Four-octet AS Number Space*
RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*
RFC 5065, *Autonomous System Confederations for BGP*
RFC 5291, *Outbound Route Filtering Capability for BGP-4*
RFC 5396, *Textual Representation of Autonomous System (AS) Numbers (asplain)*
RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6
Next Hop*
RFC 5575, *Dissemination of Flow Specification Rules*
RFC 5668, *4-Octet AS Specific BGP Extended Community*
RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*
RFC 6811, *Prefix Origin Validation*
RFC 6996, *Autonomous System (AS) Reservation for Private Use*
RFC 7311, *The Accumulated IGP Metric Attribute for BGP*
RFC 7607, *Codification of AS 0 Processing*
RFC 7674, *Clarification of the Flowspec Redirect Extended Community*
RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE)
Information Using BGP*
RFC 7854, *BGP Monitoring Protocol (BMP)*
RFC 7911, *Advertisement of Multiple Paths in BGP*
RFC 7999, *BLACKHOLE Community*
RFC 8092, *BGP Large Communities Attribute*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet
(SAToP)*
RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation
Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*

draft-ietf-bess-evpn-pref-df-01, *Preference-based EVPN DF Election*

draft-ietf-bess-evpn-prefix-advertisement-11, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-04, Operational Aspects of Proxy-ARP/ND in EVPN Networks
draft-ietf-bess-evpn-vpls-seamless-integ-03, (PBB-)EVPN Seamless Integration with (PBB-)VPLS
draft-snr-bess-pbb-evpn-isid-cmacflush-01, PBB-EVPN ISID-based CMAC-Flush
RFC 7432, BGP MPLS-Based Ethernet VPN
RFC 7623, Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)
RFC 8214, Virtual Private Wire Service Support in Ethernet VPN
RFC 8317, Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)
RFC 8365, A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)

Frame Relay

ANSI T1.617 Annex D, DSS1 - Signalling Specification For Frame Relay Bearer Service
FRF.1.2, PVC User-to-Network Interface (UNI) Implementation Agreement
FRF.12, Frame Relay Fragmentation Implementation Agreement
FRF.16.1, Multilink Frame Relay UNI/NNI Implementation Agreement
FRF.5, Frame Relay/ATM PVC Network Interworking Implementation
FRF2.2, PVC Network-to-Network Interface (NNI) Implementation Agreement
ITU-T Q.933 Annex A, Additional procedures for Permanent Virtual Connection (PVC) status management

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information
RFC 3471, Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description
RFC 3473, Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions
RFC 4204, Link Management Protocol (LMP)
RFC 4208, Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model
RFC 4872, RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery
RFC 5063, Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart (helper mode)

gRPC Remote Procedure Calls (gRPC)

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.4.0*

*gRPC Network Management Interface (gNMI), Capabilities, Get, Set, Subscribe
(ONCE, SAMPLE, ON_CHANGE)*

Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to
Intermediate system intra-domain routing information exchange protocol for
use in conjunction with the protocol for providing the connectionless-mode
Network Service (ISO 8473)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 2973, *IS-IS Mesh Groups*

RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate
System to Intermediate System*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate
System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate
System to Intermediate System (IS-IS)*

RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for
Advertising Router Information*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*

RFC 5306, *Restart Signaling for IS-IS (helper mode)*

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label
Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*
RFC 8202, *IS-IS Multi-Instance* (single topology)

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
RFC 7431, *Multicast-Only Fast Reroute*
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*
RFC 768, *User Datagram Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specifications*
RFC 1350, *The TFTP Protocol (revision 2)*
RFC 2347, *TFTP Option Extension*
RFC 2348, *TFTP Blocksize Option*
RFC 2349, *TFTP Timeout Interval and Transfer Size Options*
RFC 2428, *FTP Extensions for IPv6 and NATs*
RFC 2784, *Generic Routing Encapsulation (GRE)*
RFC 2890, *Key and Sequence Number Extensions to GRE*
RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*
RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*
RFC 4254, *The Secure Shell (SSH) Connection Protocol*
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*
RFC 5925, *The TCP Authentication Option*
RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-bier-mvpn-11, *Multicast VPN Using BIER*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

draft-ietf-mboned-mtrace-v2-17, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*

RFC 4607, *Source-Specific Multicast for IP*

RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*

RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*

RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks (MPLS encapsulation)*

RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1191, *Path MTU Discovery (router specification)*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1534, *Interoperation between DHCP and BOOTP*
RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
RFC 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 2003, *IP Encapsulation within IP*
RFC 2131, *Dynamic Host Configuration Protocol*
RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
RFC 2401, *Security Architecture for Internet Protocol*
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*
RFC 3046, *DHCP Relay Agent Information Option (Option 82)*
RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
RFC 4884, *Extended ICMP to Support Multi-Part Messages (ICMPv4 and ICMPv6 Time Exceeded)*

Internet Protocol (IP) — Version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*
RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3596, *DNS Extensions to Support IP version 6*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4193, *Unique Local IPv6 Unicast Addresses*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*
RFC 5007, *DHCPv6 Leasequery*
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5722, *Handling of Overlapping IPv6 Fragments*
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 (IPv6)*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 8201, *Path MTU Discovery for IP version 6*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*
RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
RFC 2409, *The Internet Key Exchange (IKE)*
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*
RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*
RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*
RFC 4301, *Security Architecture for the Internet Protocol*
RFC 4303, *IP Encapsulating Security Payload*
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
RFC 4308, *Cryptographic Suites for IPsec*
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*
RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*
RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*
RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*
RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*
RFC 5903, *ECP Groups for IKE and IKEv2*
RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
RFC 6379, *Suite B Cryptographic Suites for IPsec*
RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*
RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*
RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*
RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, Controlling State Advertisements Of Non-negotiated LDP Applications

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities

draft-pdutta-mpls-ldp-v2-00, LDP Version 2

draft-pdutta-mpls-mldp-up-redundancy-00, Upstream LSR Redundancy for Multipoint LDP Tunnels

draft-pdutta-mpls-multi-ldp-instance-00, Multiple LDP Instances

draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction

RFC 3037, LDP Applicability

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (helper mode)

RFC 5036, LDP Specification

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs)

RFC 5443, LDP IGP Synchronization

RFC 5561, LDP Capabilities

RFC 5919, Signaling LDP Label Advertisement Completion

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 6512, Using Multipoint LDP When the Backbone Has No Route to the Root

RFC 6826, Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 7032, LDP Downstream-on-Demand in Seamless MPLS

RFC 7552, Updates to LDP for IPv6

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions

RFC 2661, Layer Two Tunneling Protocol "L2TP"

RFC 2809, Implementation of L2TP Compulsory Tunneling via RADIUS

RFC 3438, Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update

RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3)

RFC 4719, Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)

RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"

Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 (IPv6)*

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

openconfig-bgp.yang version 3.0.1, *BGP Module*

openconfig-bgp-common.yang version 3.0.1, *BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.0.0, *Interfaces Ethernet Model*

openconfig-if-ip.yang version 2.0.0, *Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *Interfaces IP Extensions Module*

openconfig-interfaces.yang version 2.0.0, *Interfaces Module*

openconfig-isis.yang version 0.3.0, *IS-IS Module*
openconfig-isis-lsp.yang version 0.3.0, *IS-IS LSP Module*
openconfig-isis-routing.yang version 0.3.0, *IS-IS Routing Module*
openconfig-lacp.yang version 1.1.0, *LACP Module*
openconfig-lldp.yang version 0.1.0, *LLDP Module*
openconfig-local-routing.yang version 1.0.1, *Local Routing Module*
openconfig-network-instance.yang version 0.8.0, *Network Instance Module*
openconfig-routing-policy.yang version 3.0.0, *Routing Policy Module*
openconfig-vlan.yang version 2.0.0, *VLAN Module*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1212, *Concise MIB Definitions*
RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*
RFC 1215, *A Convention for Defining Traps for use with the SNMP*
RFC 1724, *RIP Version 2 MIB Extension*
RFC 1901, *Introduction to Community-based SNMPv2*
RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*
RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*
RFC 2206, *RSVP Management Information Base using SMIv2*
RFC 2213, *Integrated Services Management Information Base using SMIv2*
RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*
RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*
RFC 2515, *Definitions of Managed Objects for ATM Management*
RFC 2570, *SNMP Version 3 Framework*
RFC 2571, *An Architecture for Describing SNMP Management Frameworks*
RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
RFC 2573, *SNMP Applications*
RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*
RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*
RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
RFC 2579, *Textual Conventions for SMIv2*
RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*
RFC 4273, *Definitions of Managed Objects for BGP-4*
RFC 4292, *IP Forwarding Table MIB*
RFC 4293, *Management Information Base for the Internet Protocol (IP)*
RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*
RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*
RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*
RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
RFC 5102, *Information Model for IP Flow Information Export*
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2 (TLS client, RSA public key)*
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*
RFC 6991, *Common YANG Data Types*
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*
RFC 7950, *The YANG 1.1 Data Modeling Language*
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*
RFC 5921, *A Framework for MPLS in Transport Networks*
RFC 5960, *MPLS Transport Profile Data Plane Architecture*
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*
RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6022, *YANG Module for NETCONF Monitoring*
RFC 6241, *Network Configuration Protocol (NETCONF)*
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
RFC 6243, *With-defaults Capability for NETCONF*
RFC 7895, *YANG Module Library*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extensibility*
RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*
RFC 1765, *OSPF Database Overflow*
RFC 2328, *OSPF Version 2*
RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 4552, *Authentication/Confidentiality for OSPFv3*
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 5185, *OSPF Multi-Area Adjacency*
RFC 5187, *OSPFv3 Graceful Restart (helper mode)*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*
RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*
RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

TS-007, *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1989, *PPP Link Quality Monitoring*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*
MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*
MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*
MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*
RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*
RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*
RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*
RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*
RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*
RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*
RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*
RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*
RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*
RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*
RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*
RFC 6073, *Segmented Pseudowire*
RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
RFC 6718, *Pseudowire Redundancy*
RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
RFC 6870, *Pseudowire Preferential Forwarding Status bit*
RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)*

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-filsfils-spring-segment-routing-policy-05, *Segment Routing Policy for Traffic Engineering*

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-02, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*

draft-ietf-mpls-spring-lsp-ping-02, Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing

draft-ietf-spring-conflict-resolution-05, Segment Routing MPLS Conflict Resolution

draft-ietf-spring-segment-routing-ldp-interop-09, Segment Routing interworking with LDP

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, Jitter Network Interfaces

ANSI T1.105.06, Physical Layer Specifications

ANSI T1.105.09, Network Timing and Synchronization

ITU-T G.703, Physical/electrical characteristics of hierarchical digital interfaces

ITU-T G.707, Network node interface for the synchronous digital hierarchy (SDH)

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC)

ITU-T G.823, The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy

ITU-T G.824, The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy

ITU-T G.825, The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)

ITU-T G.841, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002

ITU-T G.957, Optical interfaces for equipments and systems relating to the synchronous digital hierarchy

Time Division Multiplexing (TDM)

ANSI T1.403, DS1 Metallic Interface Specification

ANSI T1.404, DS3 Metallic Interface Specification

Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions*, issued 09/2008

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*, issued 03/2003

ITU-T G.8261, *Timing and synchronization aspects in packet networks*, issued 04/2008

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*, issued 08/2007

ITU-T G.8264, *Distribution of timing information through packet networks*, issued 10/2008

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization*, issued 10/2010

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*, issued 07/2014

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) (server, unauthenticated mode)*

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support



Customer Documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation Feedback

[Customer Documentation Feedback](#)

