Configuration Commands

Generic Commands

description

Syntax **description** *description-string*
no description

Context config>port
config>port>ethernet>access>egr>vport
config>port>ethernet>access>egr>qgrp
config>port>ethernet>access>ing>qgrp
config>port>ethernet>network>egr>qgrp
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
config>lag>link>map>profile
config>port>ethernet>eth-cfm>mep
config>card>fp>ingress>access>queue-group
config>card>fp>ingress>network>queue-group

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 is associated with the configuration context.

Parameters *long-description-string* — 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, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>card
config>card>mda
config>interface-group-handler
config>port
config>port>ethernet
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>multilink-bundle>ima>test-pattern-procedure
config>port>sonet-sdh>path>atm>ilmi
config>lag
config>port>ethernet>eth-cfm>mep
config>port>ethernet>efm-oam
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

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.

Special Cases

card — The default state for a card is **no shutdown**.

interface group handler (IGH) — The default state for an IGH is **shutdown**.

mda — The default state for a mda is **no shutdown**.

lag — The default state for a Link Aggregation Group (LAG) is **shutdown**.

port — The default state for a port is **shutdown**.

path — The default state for a SONET/SDH path is **shutdown**.

Card Commands

card

Syntax **card** *slot-number*
no card *slot-number*

Context config

Description This mandatory command enables access to the chassis card Input/Output Control Forwarding Module (IOM), slot, MCM and MDA CLI context.

The **no** form of this command removes the card from the configuration. All associated ports, services, and MDAs must be shutdown.

Default No cards are configured.

Parameters *slot-number* — The slot number of the card in the chassis.

Values 1 — 10, depending on chassis model.

SR-7: *slot-number* = 1 — 5
SR-12: *slot-number = 1 — 10*

capability

Syntax **capability** {**sr | ess**} [**now**]

Context config>card

Description This command sets the desired capability for the associated slot and card.

By default, the capability will be set to that of the base chassis type. To set this to a non-default value, the **mixed-mode** command must be enabled at the system level.

Changing the capability of a slot or card will result in the associated slot being reset. The card-type must first be configured before the capability command can be issued.

Default **capability ess** on a 7450 chassis

Parameters **now** — This optional keyword can be added to the interactive command to force the command to be executed immediately without further question. If this keyword is not present, then the user will be presented with a question to ensure they understand that as a result of this command, the associated slots will be reset immediately to enable **mixed-mode**.

card-type

Syntax **card-type** *card-type*

no card-type

Context config>card

Description This mandatory command adds an IOM 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 MDAMCMor 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 medium severity alarm is raised. The alarm is cleared when the correct card type is installed or the configuration is modified.

A high severity alarm 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 low severity trap is issued when a card is removed that is administratively disabled.

Because the IOM-3 integrated card does 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, etc., remain in the MDA configuration context.

An appropriate alarm is raised if a partial or complete card failure is detected. The alarm is cleared when the error condition ceases.

The **no** form of this command removes the card from the configuration.

Default No cards are preconfigured for any slots.

Parameters *card-type* — The type of card to be configured and installed in that slot.

Values iom-20g, iom2-20g, iom-20g-b, iom3-20g, iom3-40g, iom3-xp, imm48-1gb-sfp, imm48-1gb-tx, imm4-10gb-xfp, imm5-10gb-xfp, imm8-10gb-xfp, imm12-10gb-SF+, imm1-40gb-tun, imm3-40gb-qsfp, imm1-oc768-tun, imm1-100g-cfp, iom3-xp, iom4-eiom3-xp, imm-2pac-fp3

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 Alcatel-Lucent 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# 2102 tmnxEqCardChipIfDownEvent
- 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, it is recommended to 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**

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.

named-pool-mode

Syntax **named-pool-mode**

Context config>card>mda
config>port

Description The named-pool-mode CLI context is used 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.

Parameters

Virtual Scheduler Commands

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 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 rates. The window size is based on 4x the minimum rate calculation interval. Any previous measured offered rates within the window are used in the averaging function.

The system separates queues into fast and slow categories and maintains a separate minimum recalc interval for each type. The default minimum recalculation times for each type are as follows:

Slow Queue

Minimum Rate Calculation Interval:	0.1875 Seconds
Averaging Window Size:	0.75 Seconds

Fast Queue

Minimum Rate Calculation Interval:	0.0625 Seconds
Averaging Window Size:	0.25 Seconds

The actual minimum rate calculation interval may be increased or decreased by using the fast-queue and/or slow-queue keywords followed by a percent value which is applied to the default interval. The default slow-queue threshold rate is 1Mbps. Once a queue is categorized as slow, its rate must rise to 1.5Mbps before being categorized as a fast queue. The categorization threshold may be modified by using the slow-queue-thresh command.

The **no** rate-calc-min-interval command is used to restore the default fast queue and slow queue minimum rate calculation interval.

Parameters **fast-queue** *percent-of-default*: — 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%

slow-queue *percent-of-default*: — 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** *percent-of-default*
no sched-run-min-int

Context config>card>virt-sched-adj

Description This command is used to override the default minimum time that must elapse before a virtual scheduler may redistribute bandwidth based on changes to the offered rates of member 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 queue offered rates is performed sequentially and the interval allows a number queues rates to be determined prior to determining the distribution of bandwidth to the 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 is used to restore the default minimum scheduler run interval for all virtual schedulers on the card.

Parameters *percent-of-default:* — 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%

task-scheduling-int

Syntax **task-scheduling-int** *percent-of-default*
no task-scheduling-int

Context config>card>virt-sched-adj

Description This command is used to override 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 is used to restore the default task scheduling interval of the card's hierarchical virtual scheduler task.

Parameters *percent-of-default:* — 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.

Values 0.01% to 1000.00%

Default 100.00%

slow-queue-thresh

Syntax **slow-queue-thresh** *kilobits-per-second*
no slow-queue-thresh

Context config>card>virt-sched-adj

Description This command is used to override the system default rate threshold where queues are placed in the "slow" queue category. Slow rate 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 queue.

The default slow queue rate is 1Mbps. The fast rate is derived by multiplying the slow rate by a factor of 1.5 resulting in a default fast rate of 1.5Mbps. The slow-queue-thresh 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** slow-queue-thresh command is used to restore the default slow queue and fast queue rate thresholds.

Parameters *kilobit-per-second:* — The kilobit-per-second parameter is required and is used to modify the default slow queue rate threshold. Defining a value of 0 forces all queues to be treated as fast rate. Defining a value of 1000 (1Mbps) returns the threshold to the default value and is equivalent to executing no slow-queue-thresh.

The fast queue rate threshold is derived by multiplying the new slow queue rate threshold by a factor of 1.5.

Values 0 to 1000000 kilobits per second

Default 1000 kilobits per second

MCM Commands

mcm

Syntax **mcm** *mcm-slot*
no mcm *mcm-slot*

Context config>card

Description This mandatory command enables access to a card's MCM CLI context to configure MCMs.

Default No MCM slots are configured by default.

Parameters *mcm-slot* — The MCM slot number to be configured. Even slot numbers 2-12 are invalid. MCM provisioning is not required to provision Compact Media Adapters (CMAs).

Values

SR-c4:	1, 3
SR-c12:	1, 3, 5, 7, 9, 11

mcm-type

Syntax **mda** *mcm-type*
no mcm *mcm-type*

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 *mcm-type* — The type of MCM to provision for that slot.

Values mcm-v1, mcm-v2

MDA Commands

mda

Syntax **mda** *mda-slot*
no mda *mda-slot*

Context config>card

Description This mandatory command enables access to a card's MDA CLI context to configure MDAs.

Default No MDA slots are configured by default.

Parameters *mda-slot* — 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 7710 SR c-12 MDA: 1, 3, 5, 7, 9, 11
7710 SR c-12 CMA: 1-12

7710 SR c-4 MDA: 1, 3
7710 SR c-4 CMA: 1-4

mda-type

Syntax **mda-type** *mda-type*
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. 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.

Note: CMAs 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.

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 the 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.

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.

Default No MDA types are configured for any slots by default.

Parameters *mda-type* — The type of MDA selected for the slot postion.

7750: m60-10/100eth-tx, m10-1gb-sfp, m16-oc12/3-sfp, m8-oc12/3-sfp, m16-oc3-sfp, m8-oc3-sfp, m4-oc48-sfp, m1-oc192, m5-1gb-sfp, m12-chds3, m1-choc12-sfp, m1-10gb, m4-choc3-sfp, m2-oc192-xp-xfp, m2-oc48-sfp, m20-100eth-sfp, m20-1gb-tx, m2-10gb-xfp, m2-oc192-xfp, m12-1gb-sfp, m12-1gb+2-10gb-xp, m4-atmoc12/3-sfp, m16-atmoc3-sfp, m20-1gb-sfp, m4-chds3, m1-10gb-xfp, vsm-cca, m5-1gb-sfp-b, m10-1gb-sfp-b, m4-choc3-as-sfp, m10-1gb+1-10gb, isa-ipsec, m1-choc12-as-sfp, m12-chds3-as, m4-chds3-as, isa-aa, isa-tms, m12-1gb-xp-sfp, m12-1gb+2-10gb-xp, m10-1gb-hs-sfp, m1-10gb-hs-xfp, m4-choc3-ces-sfp, m1-choc3-ces-sfp, m4-10gb-xp-xfp, m2-10gb-xp-xfp, m1-10gb-xp-xfp, m10-1gb-xp-sfp, m20-1gb-xp-sfp, m20-1gb-xp-tx, m1-choc12-ces-sfp, p1-100g-cfp, p10-10g-sfp,p3-40g-qsfp, p6-10g-sfp, imm24-1gb-xp-sfp, imm24-1gb-xp-tx, imm5-10gb-xp-xfp, imm4-10gb-xp-xfp, imm3-40gb-qsfp, imm1-40gb-qsfp, imm1-40gb-xp-tun, imm-1pac-fp3/p1-100g-tun, imm2-10gb-xp-xfp, imm12-10gb-xp-SF+, imm1-oc768-xp-tun, imm1-100gb-xp-cfp, isa-video, m1-10gb-dwdm-tun, iom3-xp-b, m4-atmoc12/3-sf-b, m16-atmoc3-sfp-b, m16-oc12/3-sfp-b, m4-oc48-sfp-b, me10-10gb-sfp+, me1-100gb-cfp2

7750 SR-c12/4: m60-10/100eth-tx, m8-oc3-sfp, m5-1gb-sfp, m2-oc48-sfp, m20-100eth-sfp, m20-1gb-tx, m4-atmoc12/3-sfp, m20-1gb-sfp, m5-1gb-sfp-b, m4-choc3-as-sfp, c8-10/100eth-tx, c1-1gb-sfp,c2-oc12/3-sfp-b, c8-chds1, c4-ds3, c2-oc12/3-sfp, c1-choc3-ces-sfp, m1-choc12-as-sfp, m12-chds3-as, m4-chds3-as, m4-choc3-ces-sfp, m10-1gb-xp-sfp, m20-1gb-xp-sfp, m20-1gb-xp-tx

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.

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.

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 The named-pool-policy command is used to associate 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.

Parameters *policy-name* — The defined policy-name must be an existing named pool policy on the system. If policy-name does not exist, the named-pool-policy command will fail. If another named pool policy is currently associated, it will continue 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.

Values Any existing Named Pool Policy

Default None

egress-xpl

Syntax **egress-xpl**

Context configure>card>mda

Description This command enables the context to configure **egress-xpl** settings used by the **fail-on-error** feature.

threshold

Syntax **threshold** *threshold*

Context configure>card>mda>egress-xpl

Description This command configures the Egress XPL Error Threshold value used by the **fail-on-error** feature.

Parameters *threshold* — 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 *threshold* errors per minute for *window* minutes, the MDA will be put in the *failed* state.

threshold cannot be changed while fail-on-error is enabled for this MDA.

Values 1 - 1000000

Default 1000

window

Syntax **window** *window*

Context configure>card>mda>egress-xpl

Description This command configures the Error Window value used by the fail-on-error feature.

Parameters *window* — Specifies the time (in minutes) that the MDA can experience frequent Egress XPL Errors. When **fail-on-error** is enabled, if more than *threshold* Egress XPL errors per minute occur on the MDA for <window> consecutive minutes, the MDA will be put in the *failed* state.

window cannot be changed while fail-on-error is enabled for this MDA.

Values 1 - 1440

Default 60

fail-on-error

Syntax [**no**] **fail-on-error**

Context configure>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.

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 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 — *This 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

hi-bw-mcast-src

Syntax **hi-bw-mcast-src** [**alarm***] [***group *group-id*]
no hi-bw-mcast-src

Context config>card>mda

Description This command designates the MDA 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 MDA. If a group is specified, all *MDAs 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 MDA, the MDAs will be brought online and generate an event (SYSTEM: 2052 - mdaHiBw-MulticastAlarm). 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.*

This feature is supported on the 7750 SR-7 and 7750 SR-12.

The **no** form of the command removes the high-bandwidth IP multicast source designation from the MDA.

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.

group *group-id* **—** Specifies the logical MSFP group for the MDA. MDAs configured with the same *group-id* will be placed on the same MSFP.

Values 0 — 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.

egress

Syntax **egress**

Context config>card>mda

Description This command enables the context to configure egress MDA parameters.

ingress

Syntax **ingress**

Context config>card>mda

Description This command enables the context to configure ingress MDA parameters.

mcast-path-management

Syntax **mcast-path-management**

Context config>card>mda>ingress

Description This command enables the context to configure local MDA settings for ingress multicast path management.

ancillary-override

Syntax **ancillary-override**

Context config>card>mda>ingress>mcast-mgmt

Description This command enables the context to configure ancillary path bandwidth override parameters.

path-limit

Syntax **path-limit** *megabits-per-second*
no path-limit

Context config>card>mda>ingress>mcast-mgmt>anc-override

Description This command overrides the path limits contained in the bandwidth policy associated with the MDA.

The **no** form of the command removes the path limit override from an ingress multicast path and restores the path limit defined in the bandwidth policy associated with the MDA.

Parameters *megabits-per-second* — Specifies the path limit override to give the upper limit that multicast channels may use on each path.

Values ancillary-override: 1 — 5000
primary-override: 1 — 2000
secondary-override: 1 — 2000

bandwidth-policy

Syntax **bandwidth-policy** *policy-name*
no bandwidth-policy

Context config>card>mda>ingress>mcast-mgmt

Description This command specifies an existing multicast bandwidth policy. Bandwidth policies are used to manage the ingress multicast path bandwidth. Each forwarding plane supports multicast forwarding paths into the switch fabric. Bandwidth policy parameters are configured in the **config>mcast-mgmt** context.

Parameters *policy-name* — Specifies an existing multicast bandwidth policy.

primary-override

Syntax **primary-override**

Context config>card>mda>ingress>mcast-mgmt

Description This command enables the context to configure primary path limit override parameters.

secondary-override

Syntax **secondary-override**

Context config>card>mda>ingress>mcast-mgmt

Description This command enables the context to configure secondary path limit override parameters.

scheduler-policy

Syntax **scheduler-policy** *hsmda-scheduler-policy-name*
no scheduler-policy

Context config>card>mda>ingress

Description This command overrides the default HSMDA scheduling policy on the ingress MDA. The command can only be executed on an MDA provisioned as a HSMDA. Attempting to provision a scheduler policy on a non-HSMDA will fail. The defined hsmda-scheduler-policy-name must be an existing HSMDA scheduler policy. An HSMDA scheduler policy that is currently associated with an HSMDA cannot be removed from the system.

When the scheduler policy is changed on an ingress HSMDA, the ingress scheduling parameters are immediately changed to reflect the parameters within the policy.

The scheduler policy defined on the ingress context of an HSMDA cannot be changed when local scheduler overrides exist. The scheduler overrides must be removed prior to changing the scheduler policy. Once the scheduler policy is changed, any required overrides may be redefined.

The **no** form of the command restores default HSMDA scheduler policy control over the ingress scheduler on the HSMDA. The **no scheduler-policy** command cannot be executed when scheduler overrides exist on the ingress HSMDA. The overrides must be removed prior to executing the no scheduler-policy command.

Parameters *hsmda-scheduler-policy-name* — Specifies an existing policy created in the **config>qos>hsmda-scheduler-policy** context. The "default" policy name cannot be specified. Instead, the **no scheduler-policy** command should be executed resulting in the default scheduler policy being used by the ingress MDA.

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.

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.

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.

egress

Syntax **egress**

Context config>port>access
config>card>mda>access
config>card>mda>network
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.

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** *threshold*

Context configure>card>mda>ingress-xpl

Description This command configures the Ingress XPL Error Threshold value used by the **fail-on-error** feature.

Parameters *threshold* — 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 *threshold* errors per minute for *window* minutes, the MDA will be put in the *failed* state.

threshold cannot be changed while fail-on-error is enabled for this MDA.

Values 1 - 1000000

Default 1000

window

Syntax **window** *window*

Context configure>card>mda>ingress-xpl

Description This command configures the Error Window value used by the **fail-on-error** feature.

Parameters *window* — Specifies the time (in minutes) that the MDA can experience frequent Ingress XPL Errors. When **fail-on-error** is enabled, if more than *threshold* Ingress XPL errors per minute occur on the MDA for <window> consecutive minutes, the MDA will be put in the *failed* state.

window cannot be changed while fail-on-error is enabled for this MDA.

Values 1 - 1440

Default 60

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>network>egress
config>port>network>ingress
config>port>access>uplink>egress

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 default

Parameters *name* — Specifies the pool name, a string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

resv-cbs

Syntax **resv-cbs** *percent-or-default* **amber-alarm-action step** *percent* **max** [1..100]
resv-cbs *percent-or-default*
no resv-cbs

Context config>port>access>egress>pool
config>port>ethernet>network
config>card>mda>access>egress
config>card>mda>access>ingress
config>card>mda>network>egress
config>card>mda>network>ingress
config>port>access>egress>channel>pool
config>port>access>ingress>pool
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 enqueuing 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 default (30%)

Parameters *percent-or-default* — Specifies the pool buffer size percentage.

Values 0 — 100, 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 — 100

Default 0

max [1..100] — 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 — 100

Default 0

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 0

Parameters *percentage* — Specifies the amber alarm threshold.

Values 1 — 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 0

Parameters *percentage* — Specifies the amber alarm threshold.

Values 1 — 1000

slope-policy

Syntax **slope-policy** *name*
no slope-policy

Context config>port>access>egress>pool
config>card>mda>access>egress
config>card>mda>access>ingress
config>card>mda>network>egress

config>card>mda>network>ingress
config>port>access>egress>channel>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.

General Port Commands

port

Syntax **port** {*port-id* | *bundle-id* | *bpgrp-id* | *aps-id*}
no port {*bundle-id* | *bpgrp-id* | *aps-id*}

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. (See **card** and **mda** commands.)

Default No ports are configured. All ports must be explicitly configured and enabled.

Parameters *port-id* — Specifies the physical port ID in the *slot/mda/port* format.

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The command syntax must be configured as follows:

Syntax: 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.)
bundle: keyword
slot: IOM/MDA slot numbers
bundle-num: 1 — 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 un-bundled 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 **configure>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

For example:

```
A:ALA-48>config>port>aps# info
----------------------------------------------
            working-circuit 1/2/3
            protect-circuit 4/5/6
----------------------------------------------
A:ALA-48>config>port>aps#
```

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.

Syntax: **port aps**-*group-id*
aps: keyword
group-id: 1 — 64

Example: **port aps**-64

bpgrp-id — Creates a bundle protection group (BPG). The BPGrp 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 BPGrp.

Syntax: **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.

bpg-num: 1 — 1600

Example: **port bpgrp-ima**-29

bundle-id — Specifies the multilink bundle to be associated with this IP interface. The command syntax must be configured as follows:

Syntax: bundle-*type-slot/mda.bundle-num*
bundle-ppp-*slot/mda.bundle-num* (Creates a multilink PPP bundle.)
bundle: keyword
slot: card/mda slot numbers
bundle-num: *1 — 256*

For example:
router1>config# port bundle-1/1.1 (multilink PPP bundle)

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.

channel

Syntax **channel** *channel*

Context config>port>dwdm
config>port>dwdm>tdcm
config>port>dwdm>coherent

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 (m1-10gb-dwdm-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, 17-61, 175-605]
where: 17-61 is used for 100GHz channels

175, 185 — 605 is used for 50GHz channels
0 only valid on disabled (shutdown) ports

Values The DWDM channel number range is listed in the following table.

Table 1: 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
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

Table 1: DWDM Channel Numbers (Continued)

C-Band					
100 GHz Grid			50GHz Grid		
nm	THz	ITU Channel	nm	THz	ITU Channel
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
1563.05	191.80	18	1563.45	191.75	175
1563.86	191.70	17			

cpr-window-size

Syntax **cpr-window-size** *window-size*

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] symbols

wavetracker

Syntax **wavetracker**

Context config>port>dwdm

Description This command validates whether or not the port supports Wavetracker.

Default None

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

Parameters *no power-control* — Laser output power is set to maximum.

power-control — Actively control the laser's output power to achieve the target power.

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 -20.00 dBm

Parameters *power* — Specify the desired average output power in dBm.

Values -22.00 — 3.00

target-power

Syntax **target-power** *power*

Context config>port>dwdm>coherent

Description This command configures the target transmit optical power for the port.

Default 1.00 dBm

Parameters *power* — Specify the desired average output power in dBm.

Values -20.00 — 1.00

report-alarm

Syntax **[no] report-alarm [encode-fail] [encode-degrade] [power-fail] [power-degrade] [power-high] [power-low]**

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.

Values encode-fail — Encoder failure
encode-degrade — Encoder degrade
power-fail — Power control failure
power-degrade — Power control degrade
power-high — Power control high limit reached
power-low — Power control low limit reached

encode

Syntax **encode** *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 ***wave-key*** — The *wave-key* 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 the table below:

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

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
235	3806	3820	2270	2284

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
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
485	1446	1460	1930	1944
495	1429	1443	1913	1927

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
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

dispersion

Syntax **dispersion** *dispersion*

Context config>port>dwdm>tdcm
config>port>dwdm>coherent

Description This command allows users to configure the dispersion compensation for the port when manual mode is selected.

Parameters *dispersion* — Specifies the dispersion compensation.

Values -1200—1200

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 -5000 — 5000

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.

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** [**nrdy**] [**mth**] [**mtl**] [**unlck**] [**tlim**] [**einv**] [**com**]

Context config>port>dwdm>tdcm

Description This command allows users to Enable/disable logging of tdcm alarms on the port.

Default All alarms are enabled

Parameters **nrdy** — Reports Tdcm not ready faults.

mth — 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 Tdcm module communication failure faults.

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 alarml.

rx-los-thresh

Syntax **rx-los-thresh <*threshold*>**

Context config>port>dwdm>coherent

Description This command configures the average input power LOS (Loss of Signal) threshold.

Default -23

Parameters *threshold* — Specifies port's rx los threshold.

Values -23.00 — -13.00

sweep

Syntax **sweep start** *dispersion-start* **end** *dispersion-end*

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 *dispersion-start* — Specifies the lower range limit for the dispersion compensation.

Values -1200 —1200

Default -1200

dispersion-end — Specifies the upper range limit for the dispersion compensation.

Values -1200—1200

Default 1200

sweep

Syntax **sweep start** *dispersion-start* **end** *dispersion-end*

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 *dispersion-start* — Specifies the lower range limit for the dispersion compensation.

Values -50000 — 50000

Default -25500

dispersion-end — Specifies the upper range limit for the dispersion compensation.

Values -50000 — 50000

Default 2000

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

queue-group

Syntax **queue-group** *queue-group-name* **instance** *instance-id*
no queue-group

Context config>port>ethernet>network>egress

Description This command is used to create 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 in length.

instance *instance-id* **—** Specifies the identication of a specific instance of the queue-group.

Values 1—40960

xgig

Syntax **xgig** {**lan** |**wan**}

Context config>port>ethernet

Description This command configures a 10 Gbps 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 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 the OTU encapsulation type (encapsulated 10GE-LAN/WAN or OC192). The port must be shut down before OTU is enabled.

Note: 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 **[no] fec {enhanced | g709}**

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: FEC cannot be disabled on OTU3 encapsulated OC768 or 40-Gigabit Ethernetby 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** | **11.0957**}

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 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).

11.0957 — Configures the port to transmit and receive an 11.0957 Gb/s synchronous OTU encapsulated 10GE-LAN signal (with fixed stuffing bytes in the OTU2 frame).

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 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** *threshold*

Context config>port>otu>sf-threshold

Description This command specifies the error rate at which to declare the signal fail condition for the 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 *threshold* — Specifies the signal fail (SF) threshold.

Values 3 —7

sd-threshold

Syntax **sd-threshold** *threshold*

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-*value*.

The SD threshold must::

- Be greater than the SF threshold.
- Be 5 or higher before setting the sf-sd-method to bip8.

Default 7

Parameters *threshold* — Specifies the exponent of the error rate, thus an error rate from 10E-3 to 10E-7.

Values 5 — 9

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** {**string** *string* | **bytes** *byte-sequence* | **auto-generated** | **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

string *string* — 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.

bytes — [byte1 byte2 ... byte64]. 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.

use-rx — Copies the received sm-tti to the expected either as a string or a sequence of bytes depending on the received sm-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 None

Parameters **none** — The received traffic is passed through.

squelch-rx — The received traffic is blocked.

pm-tti

Syntax **pm-tti**

Context config>port>otu

Description This command enables the context to configure path monitoring trail trace identifier parameters.

tx

Syntax **tx auto-generated**
tx *bytes bytes* [*bytes*...(up to 64 max)]
tx *string identifier*
tx auto-generated | string *identifier* **| bytes** *byte1* [*byte2*...(up to 64 bytes)]}
no tx

Context config>port>otu>pm-tti>tx

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 *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 Table 1, DWDM Channel Numbers, on page 188).

Parameters **auto-generated —** Specifies to use the system generated (default) TTI.

string *identifier* — 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.

bytes byte1 [byte2...(up to 64 bytes)] — 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.

Values 0 — FF, in hexidecimal byte notation

tx

Syntax **tx** {**auto-generated** | **string** *identifier* | **bytes** *byte1* [*byte2*...(up to 64 bytes)]}
no tx

Context config>port>otu>sm-tti>tx

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 Table 1, DWDM Channel Numbers, on page 188).

Parameters **auto-generated** — Specifies to use the system generated (default) TTI.

string *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.

bytesbyte1 [byte2...(up to 64 bytes)] — 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 — FF, in hexidecimal byte notation

tx

Syntax **tx** {*value* | **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.

value — Non-standard payload type value to transmit in the payload type field.

expected

Syntax **expected auto-generated**
expected bytes *byte* [*byte*...(up to 64 max)]
expected string *identifier*
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

string *string* — 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.

bytes — [byte1 byte2 ... 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 none, the received traffic is passed through.

Parameters **squelch-rx —** The received traffic is blocked.

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

string *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.

bytes byte1 [byte2...(up to 64 bytes)] **—** 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 — FF, in hexidecimal byte notation

expected

Syntax **expected** {**string** *string* | **bytes** *byte-sequence* | **auto-generated** | **use-rx**}

Context config>port>otu>pm-tti

Description This command allows the user to configure the expected RX 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. This trace should match the far-end port's PSI trace. When this trace does not match the received PSI trace, the OPU-TIM alarm will be reported if enabled.

Default Blank (all zeros)

Parameters **auto-generated —** Sets the default

string *string* — Sets the PSI trace to the string provided by the user. If the string is less than 64 bytes, the remaining bytes will be set to 0.

bytes — [byte1 byte2 ... 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.

use-rx — Copies the received psi-tti to the expected either as a string or a sequence of bytes depending on the received psi-tti data.

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 None

Parameters **none —** The received traffic is passed through.

squelch-rx — The received traffic is blocked.

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** {*value* | **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.

value — Expect a non-standard payload type value in the rx payload type field.

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.

Default None

Parameters **none** — The received traffic is passed through.

squelch-rx — 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**] **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**]

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-ais, otu-bdi, fec-sf, fec-sd, odu-ais, odu-oci, odu-lck, odu-bdi, opu-plm

Parameters **alarms** — Refer to the following table for alarm descriptions.

Alarm	Description
loc	Loss of lock
los	Loss of signal transitions on the data
lof	Loss of OTU framing
lom	Loss of Multi-frame
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
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.

ing-weight

Syntax **ing-weight access** *access-weight* **network** *network-weight*
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 *access-weight* — Specifies the access weight as an integer.

Values 0 to 100

Default 50

network-weight — Specifies the network weight as an integer.

Values 0 to 100

Default 50

egr-weight

Syntax **egr-weight access** *access-weight* **network** *network-weight*
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 *access-weight* — Specifies the access weight as an integer.

Values 0 to 100

Default 50

network-weight — 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.

ing-percentage-of-rate

Syntax **ing-percentage-of-rate** *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.

Parameters *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 — 1000

Default 100 (no change to active rate)

The **no** ing-percentage-of-rate command is used to remove 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%.

egr-percentage-of-rate

Syntax **egr-percentage-of-rate** *rate-percentage*
no egr-percentage-of-rate

Context config>port>modify-buffer-allocation-rate

Description The egr-percentage-of-rate command is used to increase or decrease 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.

Parameters *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)

The **no** egr-percentage-of-rate command is used to remove 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%.

egress-scheduler-override

Syntax [**no**] **egress-scheduler-override**

Context config>port>sonet-sdh>path
config>port>ethernet
config>port>tdm>ds3
config>port>tdm>ds1>channel-group
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 7750 SR OS Quality of Service Guide 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.

level

Syntax **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*]
no level *priority-level*

Context config>port>ethernet>egress-scheduler-override
config>port>sonet-sdh>path>egress-scheduler-override
config>port>tdm>ds3>egress-scheduler-override
config>port>tdm>ds1>channel-group>egress-scheduler-override
config>port>tdm>e3

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 *priority-level* — Identifies which of the eight port priority levels are being overridden.

Values 1 — 8

rate *pir-rate* **—** 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 1 — 40000000, max

cir *cir-rate* **—** 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 0— 40000000, max

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*
no max-rate

Context configure>port>ethernet>egress-scheduler-override>level>rate
configure>port>ethernet>egress-scheduler-override
configure>port>sonet-sdh>path>egress-scheduler-override>level
configure>port>sonet-sdh>path>egress-scheduler-override
configure>port>tdm>ds1>channel-group>egress-scheduler-override>level
configure>port>tdm>ds1>channel-group>egress-scheduler-override
configure>port>tdm>ds3>egress-scheduler-override>level
configure>port>tdm>ds3>egress-scheduler-override
configure>port>tdm>e1>channel-group>egress-scheduler-override
configure>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 *rate* — Specifies the explicit maximum frame based bandwidth limit. This value overrides the QoS scheduler policy rate.

Values 1 — 40000000, max

egress-scheduler-policy

Syntax **egress-scheduler-policy** *port-scheduler-policy-name*
no egress-scheduler-policy

Context config>port>ethernet
config>port>tdm>ds3
config>port>tdm>ds1>channel-group
config>port>tdm>e3
config>port>tdm>ds3>channel-group

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 queues or schedulers with port-parent associations are mapped to the appropriate port priority levels based on the port-parent command parameters. Any 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 overloaded for HSMDA. HSMDA policies should be associated with HSMDA 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 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.

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 the Ethernet LMI mode.

Default none

Parameters **none —** Specifies that theE LMI mode is set to none.

uni-n — Specifies that theE LMI mode is set to uni-n.

n393

Syntax **n393** [2..10]
no n393

Context config>port>ethernet>elmi

Description This command configures the monitored count of consecutive errors.

Parameters **2** .. **10 —** Specifies the monitored count of consecutive errors.

t391

Syntax **t391** [5..30]
no t391

Context config>port>ethernet>elmi

Description This command configures the polling timer for UNI-C.

Parameters **5** ..**30 —** Specifies the polling timer for UNI-C.

t392

Syntax **t392** [5..30]
no t392

Context config>port>ethernet>elmi

Description This command configures the polling verification timer for UNI-N.

Parameters **5** .. **30 —** Specifies the polling verification timer for UNI-N.

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, 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>interface>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>interface>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 * is used to create 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.

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.

Default **network** — Configures the Ethernet port, TDM channel or SONET path for transport network use.

access — Default channel/port mode for channelized, ASAP, and ATM MDAs.

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.

per-link-hash

Syntax **per-link-hash**
per-link-hash weighted
per-link-hash weighted auto-rebalance
no per-link-hash

Context config>lag

Description This command configured per-link-hash 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-hash on a LAG.

Parameters **weighted —** SAPs/subscribers/interfaces are distributed amongst LAG links based on SAPs/subscribers/

interfaces preconfigured class and weight. As new links are added to a LAG, existing SAPs subscribers are not impacted.

weighted auto-rebalance — SAPs/subscribers/interfaces are distributed amongst LAG links based on SAPs/subscribers/interfaces preconfigured class and weight. As new links are added to a LAG, existing SAPs are rebalanced automatically.

mac

Syntax **mac** *ieee-address*
no mac

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
config>lag
config>eth-tunnel

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 BPDU's 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 *ieee-address* — 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** *mtu-bytes*
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 or PPP-enabled port or subport 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 received larger than the MTU will be discarded. 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 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 the following table:

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 — 9212config>port>sonet-sdh>path512 — 9208
config>port>tdm>ds3 512 — 9208
config>port>tdm>ds1>channel-group512 — 9208

config>port>tdm>e3 512 — 9208
config>port>tdm>e1>channel-group512 — 9208

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** *name*
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.

queue-policy

Syntax **queue-policy** *name*
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 default

Parameters *name* — Specifies an existing network-queue policy name.

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.

Default no ppp

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, *The Point-to-Point Protocol (PPP),* 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 config>port>tdm>ds1>channel-group
config>port>tdm>ds3
config>port>tdm>e1>channel-group
config>port>tdm>e3

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 x43+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** *time-interval* [**dropcount** *drop-count*]

no keepalive

Context 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>ppp
config>port>tdm>e3>ppp

Description This command sets the keepalive interval.

The **no** form of this command returns the interval to the default value.

Default 10

Parameters *time-interval* — Specifies the time in seconds between keepalive messages, expressed as a decimal integer.

Values 1 — 60 seconds

drop-count — 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 — 255

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 Alcatel-Lucent 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.

Default none

advertise-interval

Syntax **advertise-interval** *advertise-interval*
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 *advertise-interval* — 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 — 650

hold-time

Syntax **hold-time** *hold-time*
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 *hold-time* — 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 — 650

hold-time-aps

Syntax **hold-time-aps** [**lsignal-failure** *sf-time*] [**lsignal-degrade** *sd-time*]
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 hol a specified string expression from an app-filter definition.

Default 0 (disabled)

Parameters *sf-time* — Specifies an integer to define the signal failure hold-down time in milliseconds.

Values 1 — 100

sd-time — Specifies an integer to define the signal degrade hold-down time in milliseconds.

Values 1 — 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** *ip-address*
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 *ip-address* — 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:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D

protect-circuit

Syntax **protect-circuit** *port-id*
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.

Default none

Parameters *port-id* — Specify the physical port that will act as the protection circuit for this APS group in the *slot*/*mda*/*port* format.

Syntax: *port-id*: *slot/mda/port*

Also see Modifying Hold-Down Timer Values on page 229 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 operator intervention.

Parameters *minutes* — Specify the time, in minutes, to wait before reverting back to the original working circuit after it has been restored into service.

Values 0— 60 minutes

Default 5

switching-mode

Syntax **switching-mode** {**uni-1plus1 (R8.0)| bi-directional** | **uni-directional**}

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 Signalling APS mode.

uni-directional — Configures the group to operate in Unidirectional 1+1 Signalling APS mode.

uni-1plus1 — Configures the group to operate in Unidirectional 1+1 Signalling and Datapath APS mode (7750 SR-c4/c12 platforms only).

working-circuit

Syntax **working-circuit** *port-id* [**number** *number*]

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** *port-id* 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 aps-*group-id*. 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 aps-*group-id* under the **config>port** context and its submenus and all configuration for services that use this aps-*group-id* 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.

Default none

Parameters *port-id* — Specify the physical port that will act as the working circuit for this APS group.

Syntax: *port-id*: *slot/mda/port*

numberSpecify in

Syntax: *number*: *1-2*

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 suggested to increase the default timer on the working line accordingly with the transmission delay present on the protect line. See **hold-time** on **page 310**.

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                     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                    Diag Capable         : yes
Connector Code     : LC                      Vendor OUI           : 00:01:9c
Manufacture date   : 2004/08/20 00:00:00     Media                : SONET/SDH
Serial Number      : 6331000705
Part Number        : CT2-MS1LBTD32Z2
Optical Compliance*: 00:01:00:00:00:00:00:00
Link Len 9u        : 2 kms                   Link Len Cu          : 0 m
Link Len 9u        : 20 * 100m               Link Len 62.5u       : 0 * 10m
Link Len 50u       : 0 * 10m
===============================================================================
Port Statistics
===============================================================================
                                                    Input                Output
-------------------------------------------------------------------------------
Packets                                           3870094               6656408
Discards                                                0                     0
Unknown Proto Discards                                  0
===============================================================================
A:NS044050253#
```

wtr-annexb

Syntax **wtr-annexb** *minutes*

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 *minutes* — Specify the time, in minutes, to wait to restore for Annex B mode operation

Ethernet Port Commands

ethernet

Syntax **ethernet**

Context config>port

Description This command enables access to the context to configure Ethernet port attributes.

This context can only be used when configuring Fast Ethernet, gigabit, or 10Gig 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>interface>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>interface>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 * is used to create 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 network — for Ethernet ports

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** *queue-group-name* [**instance** *instance-id*] [**create**]
no queue-group *queue-group-name* [**instance** *instance-id*]

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.

Default none

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 — 40960

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

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 **queue-group** *queue-group-name* [**instance** *instance-id*] [**create**]
no queue-group *queue-group-name*

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.

Default none

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 queue-group.

Values 1 — 40960

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 is used to control 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.

rate

Syntax **rate** {**max** | **rate**}
no rate

Context config>port>ethernet>access>egr>qgrp>agg-rate
config>port>ethernet>access>egr>vport>agg-rate
config>port>ethernet>network>egr>qgrp>agg-rate

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 etc.).

Parameters **rate** — Specifies the rate limit for the VPORT.

Values **max**, 1— 800000000, max

limit-unused-bandwidth

Syntax [**no**] **limit-unused-bandwidth**

Context 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

Description This command is used to enable (or disable) aggregate rate overrun protection on the agg-rate context.

queue-frame-based-accounting

Syntax [**no**] **queue-frame-based-accounting**

Context 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

Description This command is used to enabled (or disable) frame based accounting on all queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports.

host-match

Syntax **host-match dest** *destination-string* [**create**]
no host-match dest *destination-string*

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 **dest** *destination-string* — Specify a host match destination string up to 32 characters in length.

create — Keyword used to create the host match. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

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* [*queue-type*] [**create**]
no queue *queue-id*

Context config>port>ethernet>access>egr>qgrp>qover
config>port>ethernet>access>ing>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.

Default none

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.

Default none

Parameters **weight** *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 — 100

Default 1

cir-weight *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 — 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.

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** [**byte** | **kilobyte**]}
no burst-limit

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 The `queue burst-limit` command is used to define an explicit 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 `burst-limit` command is supported under the sap-ingress and sap-egress QoS policy queues. The command is also supported under the ingress and egress queue-group-templates queues.

The **no** form of this command is used to restore the default burst limit to the specified queue. This is equivalent to specifying burst-limit default within the QoS policies or queue group templates. When specified within a queue-override queue context, any current burst limit override for the queue will be removed and the queue's burst limit will be controlled by its defining policy or template.

Parameters **default —** The default parameter is mutually exclusive to specifying an explicit size value. When burst-limit default is executed, the queue is returned to the system default value.

size — When a numeric value is specified (size), 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 byte qualifier must be added following size.

Values 1 to 14,000 (14,000 or 14,000,000 depending on bytes or kilobytes)

Default No default for size, use the default keyword to specify default burst limit

byte — The **bytes** qualifier is used to specify that the value given for size must be interpreted as the burst limit in bytes. The byte qualifier is optional and mutually exclusive with the kilobytes qualifier.

kilobyte — The **kilobyte** qualifier is used to specify that the value given for size must be interpreted as the burst limit in Kilobytes. The kilobyte qualifier is optional and mutually exclusive with the bytes qualifier. If neither bytes nor kilobytes is specified, the default qualifier is kilobytes.

cbs

Syntax **cbs** *size-in-kbytes*
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 The cbs command is used to define 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 default

Parameters *size-in-kbytes* — 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 — 131072 or default

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

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 The **high-prio-only** command specifies the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The **no** form of this command restores the default high priority reserved size.

Parameters *percent* — The percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10.

Values 0 — 100, default

mbs

Syntax **mbs** *size-in-kbytes*
no mbs

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 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 enqueuing 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.

If the CBS value is larger than the MBS value, an error will occur, preventing the MBS change.

The **no** form of this command returns the MBS size assigned to the queue to the value.

Default default

Parameters *size-in-kbytes* — The size parameter is an integer expression of the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kbps, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 — 131072 or default

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.

The **no** form of the command removes queue depth monitoring for the specified queue.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
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 packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled as in- or out-of-profile 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*.

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 *pir-rate* — Defines the administrative PIR rate, in kilobits, 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 — 100000000, **max**

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. 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 — 100000000, **max**

Default 0

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>port>ethernet>access>egr>qgrp
config>port>ethernet>access>ing>qgrp
config>port>ethernet>network>egr>qgrp

Description This command associates a virtual scheduler policy with a port queue group. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the queue-group.

Parameters *scheduler-policy-name* — The *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers.

exp-secondary-shaper

Syntax **exp-secondary-shaper** {**default** | *secondary-shaper-name*} **create**
no exp-secondary-shaper *secondary-shaper-name*

Context config>port>ethernet>egress

Description This command configures the Ethernet egress expanded secondary shaper on this port.

Parameters *secondary-shaper-name* — Specifies the secondary shaper name to apply to this port.

default — Specifies the default secondary shaper to apply to this port.

create — Creates a new secondary shaper for this port.

rate

Syntax **rate** {**max** | *kilobits-per-second*}
no rate

Context config>port>ethernet>egress>exp-secondary-shaper

Description This command is used to configure 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 rat. 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 is used to restore the default metering and profiling rate to a policer.

Parameters {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the shaper 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.

Values 1—10000000 kbps

class

Syntax **class** *class-number* **rate** {*kilobits-per-second* | **max**} [**monitor-threshold** *size-in-kilobytes*]
no class

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-id* — Specifies the class identifier of the low burst max class for the shaper.

Values 1— 32

rate {*kilobits-per-second* | **max**} — Specifies the rate limit for the secondary shaper.

Values **max**, 1— 10000000

monitor-threshold *size-in-kilobytes* — Specifies the monitor threshold for the secondary shaper.

Values 0— 8190

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— 8

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 a port scheduler policy is applied to a Vport the following command is used:

configure>port>ethernet>acess>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 in length. 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 rate**

Context config>port>sonet-sdh>path>access>egress>vport
configure>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/port-scheduler-policy involves removing the existing command and applying the new command.

Parameters *agg-rate* — Specifies the rate limit for the Vport.

Values 1 — 800000000, max

egress-rate-modify

Syntax **[no] egress-rate-modify**

Context configure>port>ethernet>access>egress>vport
configure>port>sonet-sdh>path>access>egress>vport

Description This command is used to apply HQoS Adjustment to a Vport. HQoS Adjustment refers to the dynamic adjustment of the rate limit at an QoS enforcement point within 7x50 when the multicast traffic stream is disjointed from the unicast traffic stream. This QoS enforcement point within 7x50 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 7x50 BNG node in the form of a Vport. In case that the multicast delivery method in the 7x50 BNG utilizes redirection, the multicast traffic in the 7x50 BNG will flow outside of the subscriber or the Vport context and thus will bypass any bandwidth enforcement in 7x50. To correct this, a Vport bandwidth adjustment is necessary in 7x50 that will account for the multicast bandwidth consumption that is bypassing Vport in 7x50 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 **configure>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** *description-string* [**create**]
no host-match dest *destination-string*

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 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

mon-port-sch

Syntax **mon-port-sch**
no mon-port-sch

Context config>port>ethernet
configure>port>ethernet>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>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.

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.

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 gigabit 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 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 autonegotate 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 enabled for compliance with IEEE 801.3.

SR OS 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.

dot1q-etype

Syntax **dot1q-etype 0x0600..0xffff**
no dot1q-etype

Context config>port>ethernet

Description This command specifies the Ethertype expected when the port's encapsualtion type is dot1q. Dot1q encapsulation is supported only on Ethernet interfaces.

The **no** form of this command reverts the dot1q-etype value to the default.

Parameters *0x0600..0xffff* — Specifies the Ethertype to expect.

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 **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

discovery

Syntax **discovery**

Context config>port<port-id>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-capability

Syntax **advertise-capability**

Context config>port<port-id>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<port-id>ethernet>efm-oam>discovery>advertise-capability

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 this capability

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 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 is disabled grace is not enabled on those ports. Enabling grace during an active ISSU or soft reset will not been in for that event.

Default

config>system>ethernet>efm-oam	[no] grace-tx-enable
config>port>ethernet>efm-oam	grace-tx-enable

hold-time

Syntax **hold-time** *time-value*
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 0

Parameters *time-value* — 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.

Values 0 — 50

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, etc.) 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-monitoring>errored-frame
config>port>ethernet>efm-oam>link-monitoring>errored-frame-period
config>port>ethernet>efm-oam>link-monitoring>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** *errored-frames*
no sd-threshold

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame

Description The option is used to define 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 *errored-frames* — The number of errored frames within the configured window which indicates the port has become degraded.

Values [1… 1,000,000]

sf-threshold

Syntax **sf-threshold** *errored-frames*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame

Description The option is used to define 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 *errored-frames* — The number of errored frames within the configured window which indicates the port has become unusable.

Values [1… 1,000,000]

Default 1

window

Syntax **window** *deciseconds*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame

Description This command defines the size of the window using a 100ms base *deciseconds*. 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 *deciseconds* — The number of 100ms increments. Must be specified in increments of 10 (full seconds).

Values [10..600]

Default 10

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** *errored-frames*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame-period

Description The option is used to define 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 *errored-frames* — The number of errored frames within the configured window which indicates the port has become degraded.

Values [1… 1,000,000]

sf-threshold

Syntax **sf-threshold** *errored-frames*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame-period

Description The option is used to define 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 *errored-frames* — The number of errored frames within the configured window which indicates the port has become unusable.

Values [1… 1,000,000]

Default 1

window

Syntax **window** *packets*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame-period

Description 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 *packets* — The number of received packets.

Values [1…4,294,967,295]

Default 1,488,095 (representing 1Gbps @ 1s)

errored-frame-seconds

Syntax **errored-frame-seconds**

Context config>port>ethernet>efm-oam>link-monitoring

Description The context used to define 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** *errored-frames*
[no] sd-threshold

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame-seconds

Description The option is used to define 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** value of this option disables the sd-threshold

Default [no] sd-threshold

Parameters *errored-seconds* — The number of errored seconds within the configured window which indicates the port has become degraded.

Values [1… 900]

sf-threshold

Syntax **sf-threshold** *errored-seconds*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame-seconds

Description The option is used to define 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 *errored-seconds* — The number of errored seconds within the configured window which indicates the port has become unusable.

Values [1… 900]

Default 1

window

Syntax **window** *deciseconds*

Context config>port>ethernet>efm-oam>link-monitoring>errored-frame-seconds

Description This command defines the size of the window using a 100ms base *deciseconds*. 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 *deciseconds* — The number of 100 ms increments. Must be specified in increments of 10 (full seconds).

Values [1000..9000]

Default 600

errored-symbols

Syntax **sf-threshold** *errored-symbols*

Context config>port>ethernet>efm-oam>link-monitoring

Description The context used to define 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 **event-notification**
[no] event-notification

Context config>port>ethernet>efm-oam>link-monitoring>errored-symbols

Description This command allows the symbol error event threshold crossing actions to transmit the Event Notification OAMPDU 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 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** *errored-symbols*
[no] sd-threshold

Context config>port>ethernet>efm-oam>link-monitoring>errored-symbols

Description This option is used to define 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 *errored-symbols* — The number of errored symbols which indicates the port has become degraded.

Values [1… 1,000,000]

sf-threshold

Syntax **sf-threshold** *errored-symbols*

Context config>port>ethernet>efm-oam>link-monitoring>errored-symbols

Description The option is used to define 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 *errored-symbols* — The number of errored-symbols which indicates the port has become unusable.

Values [1… 1,000,000]

Default 1

window

Syntax **window** *deciseconds*

Context config>port>ethernet>efm-oam>link-monitoring>errored-symbols

Description Defines the size of the window using a 100ms base *deciseconds*. 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 *deciseconds* — The number of 100ms increments. Must be specified in increments of 10 (full seconds).

Values [10..600]

Default 10

shutdown

Syntax **[no] shutdown**

Context config>port>ethernet>efm-oam>link-monitoring

Description This command enables or disables 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<port-id>ethernet> efm-oam>link-monitoring>errored-frame
config>port<port-id>ethernet>efm-oam>link-monitoring>errored-frame-period
config>port<port-id>ethernet>efm-oam>link-monitoring>errored-frame-seconds
config>port<port-id>ethernet>efm-oam>link-monitoring>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 The configuration context used to define 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** *packets*

Context config>port>ethernet>efm-oam>link-monitoring>local-sf-action

Description The configuration parameters that define 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 *packets* — The number of Event Notification OAM PDUs to send to a peer when the signal failure threshold has been reached.

Values [1…5]

Default 1

info-notification

Syntax info-notification

Context config>port>ethernet>efm-oam>link-monitoring>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 OAMPDU Flag fields used to continually communicate current monitor state to the peer.

Interactions: The signal failure threshold will trigger these actions.

dying-gasp

Syntax [no] **dying-**gasp

Context config>port>ethernet>efm-oam>link-monitoring>local-sf-action>info-notification

Description The configuration option will set the dying gasp Flag field in the Information OAMPDU when the local signal failure (sf-threshold) threshold is reached. This will be maintained in all subsequent Information OAMPDUs until the situation is cleared.

Interactions: The signal failure threshold will trigger these actions.

Default no dying-gasp

critical-event

Syntax [**no**] **critical-event**

Context config>port>ethernet>efm-oam>link-monitoring>local-sf-action>info-notification

Description The configuration option will set 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 OAMPDUs until the situation is cleared.

Interactions: The signal failure threshold will trigger these actions.

Default no critical-event

local-port-action

Syntax **local-port-action** {**log-only** | **out-of-service**}

Context config>port>ethernet>efm-oam>link-monitoring>local-sf-action

Description The configuration 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 activites (such as loopback) with the peer.

Default 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 a 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 faul 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** *interval* [**multiplier** *multiplier*]

Context config>port>ethernet>efm-oam

Description This command configures the transmit interval of OAM PDUs.

Default transmit-interval 10 multiplier 5

Parameters *interval* — Specifies the transmit interval.

Values 1 — 600 (in 100 milliseconds)

multiplier *multiplier* **—** Specifies the multiplier for transmit-interval to set local link down timer.

Values 2 — 5

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 **no** form of this command returns the value to the default.

Default no egress-rate

Parameters *sub-rate* — The egress rate in Kbps.

Values 1 — 10000000

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 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 con-

figured on a port with a null encapsulation type.

qinq — Specifies QinQ encapsulation.

hold-time

Syntax **hold-time** {[**up** *hold-time up*] [**down** *hold-time down*] [**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 is used to dampen 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 is used to synchronizing 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 **up** *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 — 36000 seconds
0, 10 — 3600000 centiseconds in 5 centisecond increments

down *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 — 36000 seconds
0, 10 — 3600000 centiseconds in 5 centisecond increments

seconds | **centiseconds** — Specify the units of your hold time in **seconds** or **centiseconds**.

Note:

hsmda-scheduler-overrides

Syntax [**no**] **hsmda-scheduler-overrides**

Context config>port>ethernet

Description This command enables the context to configure ingress and egress HSMDA scheduler override parameters. Executing hsmda-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:

Values	Command	Configuration
	description	no description
	max-rate	no max-rate
	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 HSMDA egress port or ingress MDA scheduler. Once existing overrides are removed, the scheduler reverts all scheduling parameters back to the parameters defined on the hsmda-scheduler-policy associated with the egress port or ingress MDA.

group

Syntax **group** *group-id* **rate** *rate*
no group *group-id*

Context config>port>ethernet>hsmda

Description This command changes the maximum rate allowed for a weighted scheduling group on the local HSMDA scheduler. Scheduling classes within the group are managed with an aggregate rate limit when either an explicit group rate is defined on the HSMDA 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 HSMDA 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 HSMDA 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 HSMDA scheduler policy for the weighted scheduling group allowing it to use as much total bandwidth as possible.

Values 1 — 40000, max (Mbps)

max-rate

Syntax **max-rate** *rate*
no max-rate

Context config>port>ethernet>hsmda

Description This command overrides the **max-rate** parameters configured in the hsmda-scheduler-policy associated with the egress port or ingress MDA. When a **max-rate** is defined at the override level, the HSMDA scheduler policy's **max-rate** parameter is ignored.

The **hsmda-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 HSMDA 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 HSMDA scheduler policy associated with the port or MDA will be used by the local scheduler context.

Parameters *rate* — 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 — 40000000, max (Mbps)

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 HSMDA scheduler context.

scheduling-class

Syntax **scheduling-class** *class* **rate** *rate*
scheduling-class *class* **weight** *weight-in-group*
no scheduling-class *class*

Context config>port>ethernet>hsmda

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 keyword max 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 *class* — Identifies the scheduling class to be being overridden.

Values 1 — 8

rate — 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 — 40000000, max (Mbps)

weight *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 — 100

ingress-rate

Syntax **ingress-rate** *sub-rate*
no ingress-rate

Context config>port>ethernet

Description This command configures the maximum amount of ingress bandwidth that this port can receive.

The ingress-rate command is only valid for oversubscribed Ethernet MDAs. See Oversubscribed Ethernet MDAs on page 25 for details.

The **no** form of this command returns the value to the default.

Default no ingress-rate

Parameters *sub-rate* — The egress rate in mbps.

Values 1 — 10000 mbps

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.

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** *option*
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 on page 128). 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.

pbb-etype

Syntax **pbb-etype** [**0x0600**..**0xffff**]
no pbb-etype

Context config>port>ethernet

Default 0x88E7

Description This command configures the Ethertype used for PBB encapsulation.

Values **0x0600**..**0xffff:** 1536 — 65535 (accepted in decimal or hex)

qinq-etype

Syntax **qinq-etype** *0x0600..0xffff*
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.

Parameters *0x0600..0xffff* — Specifies the qinq-etype to expect.

Values 1536 — 65535 in decimal or hex formats.

report-alarm

Syntax [**no**] **report-alarm** [**signal-fail**] [**remote**] [**local**] [**no-frame-lock**] [**lcd**]

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.

lcd — Reports a codegroup delineation error.

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**}

Context config>port>ethernet

Description This command configures the port speed of a Fast Ethernet port when autonegotiation is disabled. 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 **100**

Parameters **10 —** Sets the link to 10 mbps speed.

100 — Sets the link to 100 mbps speed.

1000 — Sets the link to 1000 mbps speed.

ssm

Syntax **ssm**

Context config>port>ethernet

Description This command enables Ethernet Synchronous Status Message (SSM).

code-type

Syntax **code-type** [**sonet | sdh**]

Context config>port>ethernet>ssm

Description This command configures the encoding of synchronous 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). For compatibility with Release 7.0, sdh is the default.

Default 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*10E-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 — 9

multiplier *multiplier* **—** Specifies the multiplier used to scale the symbol error ratio.

Values 1 — 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*10E-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 — 9

multiplier *multiplier* **—** Specifies the multiplier used to scale the symbol error ratio.

Values 1 — 9

window-size

Syntax **window-size** *seconds*
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 10

Parameters *seconds* — Specifies the size of the sliding window in seconds over which the errors are measured.

Values 5 — 60

xgig

Syntax **xgig** {**lan** |**wan**}

Context config>port>ethernet

Description This command configures a 10 Gbps 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 **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.

Default none

sd-threshold

Syntax **sd-threshold** *threshold* [**multiplier** *multiplier*]
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*10E-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 **value** *threshold* **—** Specifies specifies the threshold value.

Values 1 — 9

value *multiplier* **—** Specifies specifies the multiplier value.

Values 1 — 9

sf-threshold

Syntax **sf-threshold** *threshold* [**multiplier** *multiplier*]
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*10E-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 **value** *threshold* **—** Specifies specifies the threshold value.

Values 1 — 9

value *multiplier* **—** Specifies specifies the multiplier value.

Values 1 — 9

window-size

Syntax **window-size** *seconds*
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.

Default 10

Parameters **value W —** The size of the sliding window in seconds over which the errors are measured.

Values 1-10

down-on-internal-error

Syntax [**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.

Default no down-on-internal-error

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

802.1x Port Commands

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 7750 SR 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 2

Parameters *max-auth-request* — The maximum number of RADIUS retries.

Values 1 — 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 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 7750 SR and the host can initiate an authentication procedure. The port will remain in un-authorized 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** *seconds*
no quiet-period

Context config>port>ethernet>dot1x

Description This command configures the period between two authentication sessions during which no EAPOL frames are sent by the 7750 SR.

The **no** form of this command returns the value to the default.

Default 30

Parameters *seconds* — Specifies the quiet period in seconds.

Values 1 — 3600

radius-plcy

Syntax **radius-plcy** *name*
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 *name* — Specifies an existing 802.1x RADIUS policy name.

re-auth-period

Syntax **re-auth-period** *seconds*
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 3600

Parameters *seconds* — The re-authentication delay period in seconds.

Values 1 — 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 7750 SR 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 re-authentication

server-timeout

Syntax **server-timeout** *seconds*
no server-timeout

Context config>port>ethernet>dot1x

Description This command configures the period during which the 7750 SR waits for the RADIUS server to responds to its access request message. When this timer expires, the 7750 SR 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 30

Parameters *seconds* — The server timeout period in seconds.

Values 1 — 300

supplicant-timeout

Syntax **supplicant-timeout** *seconds*
no supplicant-timeout

Context config>port>ethernet>dot1x

Description This command configures the period during which the 7750 SR 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 30

Parameters *seconds* — The server timeout period in seconds.

Values 1 — 300

transmit-period

Syntax **transmit-period** *seconds*
no transmit-period

Context config>port>ethernet>dot1x

Description This command configures the period after which the 7750 SR sends a new EAPOL request message.

The **no** form of this command returns the value to the default.

Default 30

Parameters *seconds* — The server transmit period in seconds.

Values 1 — 300

tunneling

Syntax **tunneling**
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 10Gig EthernetFast Ethernet, gigabit or 10Gig EthernetFast Ethernet or gigabit 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 — 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 *timer* — Specifies the minimum wait time before re-enabling port after loop detection.

Values 0, 10 — 160

use-broadcast-address

Syntax [**no**] **use-broadcast-address**

Context config>port>ethernet>dwl

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.

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** {*bridge-mac*}

Context config>port>ethernet>lldp

Description This command configures destination MAC address parameters.

Parameters **bridge-mac —** 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.

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 transmitw 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.

portid-subtype

Syntax **portid-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 SAM 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 SAM's ability to build those Layer Two topologies.

Default portid-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 (ie 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.

tx-mgmt-address

Syntax **tx-mgmt-address** [**system**] [**system-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 IP Address, the system IPv6 address or both. Note the system address will only be sent once. When

both options are configured both system addresses are sent. The system address must be configured for the specific version of the protocol in order to sent 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.

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.

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.

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** *policy-id*
no accounting-policy

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>sonet-sdh>path>network
config>port>tdm>ds1>network
config>port>tdm>ds3>network
config>port>tdm>e1>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 acccounting 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 — 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>ds3>network
config>port>tdm>e1>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 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>network
config>port>tdm>ds1>network
config>port>tdm>ds3>network
config>port>tdm>e1>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 default

Parameters *name* — Specifies an exisiting network-queue policy name.

Interface Group Handler Commands

interface-group-handler

Syntax [**no**] **interface-group-handler** *group-id*

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.

Default None

Parameters *group-id* — Identifies the specific Interface Group Handler.

Values 1—100

member

Syntax [**no**] **member** *portid*

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.

Default None

Parameters *portid* — Identifies the port to be associated with the interface group handler.

threshold

Syntax **threshold** *min*
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: 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 None

Parameters *min* — Specifies the minimum number of active links that must be present for the interface group handler to be active.

Values 1 — 8

Multilink-Bundle Port Commands

multilink-bundle

Syntax [**no**] **multilink-bundle**

Context config>port

Description This command creates the context to configure bundle properties for this bundle port.

Default None

fragment-threshold

Syntax **fragment-threshold** *fragment-threshold*
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 128

Parameters *fragment-threshold* — 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 — 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** *port-id*

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.

Default None

Parameters *port-id* — Specifies the physical port ID.

Syntax: *slot/mda/port.channel*

minimum-links

Syntax **minimum-links** *minimum-links*
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 1

Parameters *minimum-link* — Specify the minimum link limit, expressed as an integer.

Values 1 — 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 [**no**] **identifier** *bundle-id-string*

Context config>port>multilink-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 *bundle-id string* — Specifies the bundle ID string.

identifier

Syntax [**no**] **identifier** *frf16-link-id-string*

Context config>port>tdm>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 *frf16-link-id-string* — Specifies the bundle ID string.

ingress

Syntax **ingress**

Context config>port>multilink-bundle>mlfr

Description This command enables the context to configure the ingress QoS profile for the MLFR bundle.

egress

Syntax **egress**

Context config>port>multilink-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** *profile-id*
no qos-profile

Context config>port>multilink-bundle>mlfr>ingress
config>port>multilink-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 *profile-id* — Specifies the profile number. The value can only be modified if the MLFR bundle or FR port is shut down.

Values 1-128

frame-relay

Syntax **frame-relay**

Context config>port>multilink-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>multilink-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** *intervals*
no n391dte

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command confgures the number of DTE full status polling intervals for the LMI.

Parameters *ntervals* — 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 — 255

n392dce

Syntax **n392dce** *threshold*
no n392dce

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command configures the DCE error threshold for the LMI.

Default 3

Parameters *threshold* — Specify the number of errors that will place the bundle in an operationally down state.

Values 1 — 10

n392dte

Syntax **n392dte** *threshold*
no n392dte

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command configures the DTE error threshold for the LMI.

Parameters *count* — Specify the number of errors that will place the bundle in an operationally down state.

Values 1 — 10

n393dce

Syntax **n393dce** *count*
no n393dce

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command configures the DCE monitored event count for the LMI.

Parameters *count* — Specify the diagnostic window used to verify link integrity on the DCE interface.

Values 1 — 10

n393dte

Syntax **n393dte** *count*
no n393dte

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command configures the DTE monitored event count for the LMI.

Parameters *count* — Specify the diagnostic window used to verify link integrity on the DTE interface.

Values 1 — 10

t391dte

Syntax **t391dte** *keepalive*
no t391dte

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command configures the DTE keepalive timer value for the LMI.

Parameters *keepalive* — Specify the interval in seconds between status inquiries issued by the DTE.

Values 5 — 30

t392dce

Syntax **t392dce** *keepalive*
no t392dce

Context config>port>multilink-bundle>mlfr>frame-relay

Description This command configures the DCE keepalive timer value for the LMI.

Parameters *keepalive* — Specify the expected interval in seconds between status inquiries issued by the DTE equipment.

Values 5 — 30

hello-timeout

Syntax **hello-timeout** *seconds*
no hello-timeout

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 10 seconds

Parameters *seconds* — [1-180 seconds]

ack-timeout

Syntax **ack-timeout** *seconds*
no ack-timeout

Context config>port>multilink-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 4 seconds

Parameters *seconds* — [1-10 seconds]

retry-limit

Syntax **retry-limit** *integer*
no retry-limit

Context config>port>multilink-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 2

Parameters *integer* — Specifies the number of retransmission attempts.

Values 1 — 5

frf.12

Syntax **frf.12**

Context config>port>tdm>ds1>channel-group>frame-relay
config>port>tdm>e1>channel-group>frame-relay
config>port>tdm>ds3>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.

fragment-threshold

Syntax **fragment-threshold** *fragment-threshold*
no fragment-threshold

Context 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 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 128

Parameters *fragment-threshold* — Specifies the maximum fragment length, in bytes, to be transmitted across the FRF.12 port.

Values 128-512 bytes

mlppp

Syntax **mlppp**

Context config>port>ml-bundle

Description This command enables the context to configure multi-link PPP bundle attributes.

egress

Context config>port>ml-bundle>mlppp

Description This command enables the context to configure egress MLPPP QoS profile parameters for the multilink bundle.

Default none

ingress

Context config>port>ml-bundle>mlppp

Description This command enables the context to configure ingress MLPPP QoS profile parameters for the multilink bundle.

Default none

qos-profile

Syntax **qos-profile** *profile-id*
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.

profile-id — 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 — 128

endpoint-discriminator

Syntax **endpoint-discriminator class** {**ip-address** | **global-mac-address | null**} [**discriminator-id** *discriminator-id*]
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 descriminator 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 *discriminator-id* **—** Specifies the endpoint discriminator identifier value within the specified endpoint-discriminator class. The values below are only for the ip-address class.

Values Any valid IP address.

qos-profile

Syntax **qos-profile** *profile-id*
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 *profile-id* — 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 — 128

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** *count*
no multiclass

Context config>port>ml-bundle>multiclass

Description This command enables multi-class MLPPP as defined by RFC 2686*, The Multi-Class Extension to Multi-Link PPP*, 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 indices 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 4

Parameters *count* — Specifies the number of classes in a MLPPP bundle.

Values 2 — 4

stateless-aps-switchover

Syntax [**no**] **stateless-aps-switchover**

Context config>port>multilink-bundle> mlppp

Description This command specifies whether the bundle will perform a statefull 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 statefull switchover implies that after an APS switchover the PPP state of the bundle will be restored based on the bpgrp bundle state before the switchover.

The state cannot be changed for normal MLPPP bundles (only applicable for bpgrps). The following message appears:

The **no** form of the command disables stateless APS switchover.

Default disabled

mrru

Syntax **mrru** *mrru*
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 1524

Parameters *bytes* — Specify the maximum received reconstructed unit size, expressed as an integer.

Values 1500 — 9206 bytes

protect-bundle

Syntax [**no**] **protect-bundle** *bundle-id*

Context config>port>multilink-bundle

Description This command configures a protect bundle that is part of this BPGrp.

Parameters *bundle-id* — Specifies the protection multilink bundle in the bundle protection group. The command syntax must be configured as follows:

Syntax: bundle-*type-slot/mda.bundle-num*
bundle-PPP or IMA-*slot/mda.bundle-num* (Creates an MLPPP or IMA bundle.)
bundle: keyword
slot: IOM/MDA slot numbers
bundle-num: *1 — 256*

For example:
router1>config>port>ml-bundle> protect-bundle bundle-ima-1/1.1

red-differential-delay

Syntax **red-differential-delay** *red-diff-delay* [**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 unidirectional differential delay for IMA bundles.

The **no** form of this command restores the red-differential-delay defaults.

Default None

Parameters *red-diff-delay* — Specify the maximum red differential delay value.

Values 0 — 25 milliseconds for all other bundles
0 — 50 milliseconds for IMA 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** *bundle-id*

Context config>port>multilink-bundle

Description This command configures a working bundle that is part of this BPGrp.

Parameters *bundle-id* — Specifies the working multilink bundle in the bundle protection group. The command syntax must be configured as follows:

Syntax: bundle-*type-slot/mda.bundle-num*
bundle-PPP or IMA-*slot/mda.bundle-num* (Creates an MLPPP or IMA bundle.)
bundle: keyword
slot: IOM/MDA slot numbers
bundle-num: *1 — 256*

For example:
router1>config>port>ml-bundle> working-bundle bundle-ima-1/1.1

yellow-differential-delay

Syntax **yellow-differential-delay** *yellow-diff-delay*
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.

The **no** form of this command removes the yellow-differential-delay.

Default None

Parameters *yellow-diff-delay* — Specify the maximum yellow differential delay threshold value.

Values 1— 25 milliseconds

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**} *milli-seconds*
no link-delay {**activate** | **deactivate**}

Context config>port>multilink-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** *milli-seconds* — 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 — 30000 milli-seconds

Default 10000

deactivate *milli-seconds* — 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 — 30000 milli-seconds

Default 2000

max-bandwidth

Syntax **max-bandwidth** *number-links*
no max-bandwidth

Context config>port>ml-bundle>ima

Description This command specifies the number of links that is 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) =

(number-links) * (M-1)/M * (2048/2049) * primary member link speed

where,

M is the IMA frame size (128)
primary member link speed is either E-1 — 1920kbps or DS-1 — 1539kbps. 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 8

Parameters *number-links* — Specifies the number of links that is used to determine the maximum configurable band-

width that is allowed to be used for this IMA group.

Values 1 — 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** *port-id*
no test-link

Context config>port>ml-bundle>ima>test-pattern-procedure

Description This comand 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 *port-id* — The port ID to be used to verify link connectivity within an IMA group.

Values

port-id	slot/mda/port[.channel]
aps-id	aps-*group-id*[.*channel*]
aps	keyword
group-id	1 — 64

test-pattern

Syntax **test-pattern** *pattern*
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 0

Parameters *pattern* — Specifies an integer taking the following values:

Values 0 — 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** *IMA-version*
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 1-1

Parameters *IMA-version* — Specifies the IMA version for this group.

Values 1-0 — IMA version 1-0
1-1 — IMA version 1-1

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.

The 10 Gigabit Ethernet LAN port also has SONET/SDH characteristics. However, these characteristics are predetermined and not configurable.

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.

Note: 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 following table show MDAs that support loop timing:

Sonet/SDH	Loop Timed	Default
OC-768	Yes	node-timed
OC-192	Yes	loop-timed
OC-48	Yes	loop-timed
OC-12	No	node-timed
OC-3	No	node-timed
Channelized OC-12	Yes	loop-timed
Channelized OC-3	Yes	loop-timed
Channelized ASAP OC-12	Yes	loop-timed
Channelized ASAP OC-3	Yes	loop-timed
CES OC-3	Yes	loop-timed
ATM OC-12	No	node-timed
ATM OC-3	No	node-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.

framing

Syntax **framing** {**sonet** | **sdh**}

Context config>port>sonet-sdh

Description This command specifies SONET/SDH framing to be either SONET or SDH.

Default sonet

Parameters **sonet —** Configures the port for SONET framing.

sdh — Configures the port for SDH framing.

group

Syntax **group** *sonet-sdh-index* **payload** {**tu3** | **vt2** | **vt15**}

Context config>port>sonet-sdh

Description This command configures payload of the SONET/SDH group.

For example:

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
```

Default none

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.

tu3 — Specify 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** *hold-time* {[**up** *hold-time up*] [**down** *hold-time down*]}
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.

Default no hold-time

Parameters **up** *hold-time up* — Configures the hold-timer for link up event dampening. A value of zero (0) indicates that an up transition is reported immediately.

Values 0 — 100 in 100s of milliseconds

down *hold-time* ***down*** — The hold-timer for link down event dampening. A value of zero (0) indicates that a down transition is reported immediately.

Values 0 — 100 in 100s of milliseconds

Note: For APS configurations, the **hold-time down** and **up** default values are 100 ms and 500 ms respectively. But, 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 **aps** on page 224.

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.

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

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** *value* | **string** *string*}

Context config>port>sonet-sdh

Description This command configures the section trace bytes in the SONET section header to interoperate 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.

Default byte *0x1*

Parameters *increment-z0* — Configure an incrementing STM ID instead of a static value.

byte *value* — Set values in SONET header bytes.

Default 0x1

Values 0 — 255 or 0x00 — 0xFF

string *string* — Specifies a text string that identifies the section.

Values A string up to 16 bytes.

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 etc.

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.

Default oc12

Parameters **oc3 —** set the speed of the port to OC-3.

oc12 — Set 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.

Default no tx-dus

threshold

Syntax **threshold** {**ber-sd** | **ber-sf**} **rate** *threshold-rate*
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 line & section command), furthermore if the failure threshold is crossed the link will be set to operationally down.

Note: 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. See port aps-id on page 184.

The **no** form of this command reverts to the default value.

Default **threshold ber-sf 6** — Signal degrade BER threshold of 10^{-6}

threshold ber-sf 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

threshold-rate — The BER negative exponent (n in 10^{-n}), expressed as a decimal integer.

Values 3 — 9 (10^{-3} — 10^{-9})

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.

Default full channel (or clear channel)

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.

Syntax: sts1-*x.x*

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**}

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.

Default none

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.

report-alarm

Syntax [**no**] **report-alarms** [**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.

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 codegroup 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 16 for OC-3, DS-1, DS-3
32 for OC-12, OC-48, ATM-OC12/3, ATMOC-3, etc.

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 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 mplscp 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.

keepalive

Syntax **keepalive** *time-interval* [**dropcount** *count*]
no keepalive

Context config>port>sonet-sdh>path>ppp

Description This command enables the sending of keepalive messages and configures the time between messages and how many reports can be missed before bringing the link down.

The **no** form of this command disables the sending of echo requests.

Default **keepalive 10 dropcount 3**

Parameters *time-interval* — The time interval, in seconds, that echo requests are issued.

Values 1 — 60

Default 10

dropcount *count* **—** The number of keepalive messages that can be missed before the line is brought down.

Values 1— 255

Default 3

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 x43+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*

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.

Default 0xcf

Parameters *value* — Specifies the C2 byte value, expressed as a decimal integer or a value in hex format.

Values 1 — 254 or 0x01 — 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.

Default The default J1 value is Alcatel-Lucent XXX YYY (for example, Alcatel 7750 SR) where XXX is the platform name, such as "7750", and YYY is the product name, 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.

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.

Default 10

Parameters *time-interval* — Specifies the interval used to send periodic keepalive packets.

Values 0 — 300 seconds. A value of 0 means no keepalive packets are sent.

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 1

Parameters *up-count* — Specifies the number of continual keepalive packets that must be received in order to declare the link up.

Values 1 — 3

ATM Interface Commands

atm

Syntax **atm**

Context config>port>sonet-sdh>path
config>port>tdm>ds1>channel-group
config>port>tdm>ds3
config>port>tdm>e3
config>port>tdm>e1>channel-group
config>port>multilink-bundle>ima

Description This command enables the context to configure ATM interface properties.

cell-format

Syntax **cell-format** *cell-format*

Context config>port>tdm>ds1>channel-group>atm
config>port>tdm>ds3>atm
config>port>tdm>e3>atm
config>port>tdm>e1>channel-group>atm
config>port>multilink-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** *mapping*

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 direct cell mapping

Parameters *mapping* — 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 maping.

min-vp-vpi

Syntax **min-vp-vpi** *value*

Context config>port>sonet-sdh>path>atm
config>port>multilink-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 *value* — Specify the minimum allowable VPI value that can be used on the ATM interface for a VPC.

Values 0 — 4095 (NNI)
0 — 255 (UNI)

Default 0

ilmi

Syntax **ilmi** [*vpi/vci*]
no ilmi

Context config>port>sonet-sdh>path>atm

Description This command creates an ILMI link PVCC by default on VPI/VCI 0/16. Deleting an ILMI link deletes the PVCC. ILMI is supported only on ATM interfaces on SONET/SDH paths.

Parameters *vpi/vci* — Specifies the PVC identifier (vpi/vci).

Values

vpi	0 — 4095 (NNI) 0 — 255 (UNI)
vci	1, 2, 5 — 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** *traffic-desc-profile-id*
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 atm-td-profile 1

Parameters *traffic-desc-profile-id* — Specifies an existing ATM traffic descriptor profile. Traffic descriptor profiles are configured in the **config>qos>atm-td-profile** context.

Values 1 — 1000

keep-alive

Syntax **keep-alive** [**poll-frequency** *seconds*] [**poll-count** *value*] [**test-frequency** *seconds*]
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 devault 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** *seconds* — 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 — 255

poll-count *value* — 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 — 255

test-frequency *seconds* — Specifies the frequency for testing for connectivity when the link is establishing

before polling begins.

Values 0 — 255

protocol

Syntax **protocol** *protocol-type*
no protocol

Context config>port>sonet-sdh>path>atm>ilmi

Description This command configures the protocol.

Parameters *protocol-type* — The *protocol-type* is an enumerated integer whose value indicates the ILMI version of either 3.1 or 4.0 that is advertised by IME and also indicates the ILMI IME type of either user-side or network-side.

Values 4_0-user, 4_0-network. 3_1-user, 3_1-network

Frame Relay Commands

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>e1>channel-group>frame-relay
config>port>tdm>ds3>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>multilink-bundle>mlfr
config>port>tdm>ds1>channel-group>frame-relay
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** *profile-id*
no qos-profile

Context 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 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 *profile-id* — Specifies the profile number. The value can only be modified if the FR port is shut down.

Values 1-128

fragment-threshold

Syntax **fragment-threshold** *fragment-threshold*
no fragment-threshold

Context config>port>tdm>channel-group>frame-relay>frf.12
config>port>sonet-sdh>path>frame-relay>frf.12

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 128

Parameters *fragment-threshold* — Specifies the maximum fragment length, in bytes, to be transmitted across the FRF.12 port.

Values 128-512 bytes

identifier

Syntax [**no**] **identifier** *frf16-link-id-string*

Context config>port>tdm>ds1>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 *frf16-link--id-string* — Specifies the bundle ID string.

Values 50 chars maximum

lmi-type

Syntax **lmi-type** {**ansi** | **itu** | **none** | **rev1**}
no lmi-type

Context 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** — Use ANSI T1.617 Annex D.

itu — Use ITU-T Q933 Annex A.

none — Disable Frame Relay LMI on the given port/channel.

rev1 — Use the Rev 1 version of ANSI T1.617 Annex D.

mode

Syntax **mode** {**dce** | **dte** | **bidir**}

Context 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 enquiries over the interface. The DCE mode causes the router to respond to status enquiries. 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 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>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 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 — 255

n392dce

Syntax **n392dce** *threshold*
no n392dce

Context 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 3

Parameters *threshold* — Specify the number of errors that will place the channel in an operationally down state.

Values 1 — 10

n392dte

Syntax **n392dte** *count*
no n392dte

Context 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 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 3

Parameters *count* — Specify the number of errors that will place the path or channel in an operationally down state.

Values 1 — 10

n393dce

Syntax **n393dce** *count*
no n393dce

Context 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 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 4

Parameters *count* — Specify the diagnostic window used to verify link integrity on the DCE interface.

Values 1 — 10

n393dte

Syntax **n393dte** *number*
no n393dte

Context 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 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 4

Parameters *number* — Specify the diagnostic window used to verify link integrity on the DTE interface.

Values 1 — 10

t391dte

Syntax **t391dte** *seconds*
no t391dte

Context 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 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 10

Parameters *seconds* — Specify the interval in seconds between status inquiries issued by the DTE.

Values 5 — 30

t392dce

Syntax **t392dce** *seconds*
no t392dce

Context 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 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 15

Parameters *seconds* — Specify the expected interval in seconds between status inquiries issued by the DTE equipment.

Values 5 — 30

TDM Commands

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.

Default None

ds1

Syntax [**no**] **ds1** *ds1-id*

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 Mbps) 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 Mbps through the network. The corresponding European carrier is E-1 with a data rate of 2.048 Mbps. E-1 and T-1 (DS-1) can be interconnected for international use.

The **no** form of this command disables DS-1 capabilities.

Default None

Parameters *ds1-id* — Identifies the DS-1 channel being created.

Values DS1: 1 — 28

ds3

Syntax [**no**] **ds3** [*sonet-sdh-index*]

Context config>port>tdm

Description This command enables the context to configure DS-3 parameters. DS-3 lines provide a speed of 44.736 Mbps and is also frequently used by service providers. DS-3 lines carry 28 DS-1 signals and a 44.736 Mbps data rate.

A DS-3 connection typically supports data rates of about 43 Mbps. A T-3 line actually consists of 672 individual channels, each supporting 64 Kbps. 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.

Default none

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.

e1

Syntax **e1** [*e1-id*]

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 Mbps.

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.

Default none

Parameters *e1-id* — Specifies the E-1 channel being created.

Values E1: 1 — 21

e3

Syntax **e3** *sonet-sdh-index*

Context config>port>tdm

Description This command enables the context to configure E-3 parameters. E-3 lines provide a speed of 44.736 Mbps and is also frequently used by service providers. E-3 lines carry 16 E-1 signals with a data rate of 34.368 Mbps.

A E-3 connection typically supports data rates of about 43 Mbps. A E-3 line actually consists of 672 individual channels, each supporting 64 Kbps. 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.

bert

Syntax **bert** {**2e3**|**2e9**|**2e11**|**2e15**|**2e20**|**2e20q**|**2e23**|**ones**|**zeros**|**alternating**} **duration** *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 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 zeroes 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** *rate*
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: This command is not saved in the router configuration between boots.

Default no bit-error-insertion

Parameters *rate* — Specifies the bit error rate, expressed as an integer.

Values 2 — 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 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 hold-time** {**[up hold-time up]** [**down** *hold-time* **down]**}
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.

Default no hold-time

Parameters **up** *hold-time* **up** — Configures the hold-timer for link up event dampening. A value of zero (0) indicates that an up transition is reported immediately.

Values 0 — 100 in 100s of milliseconds (default 0)

down *hold-time* **down** — The hold-timer for link down event dampening. A value of zero (0) indicates that a down transition is reported immediately.

Values 0 — 100 in 100s of milliseconds (default 5)

This command is only supported on the m4-chds3-as, m12-chds3-as, and c4-ds3 MDAs.

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:

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 — 133	For line length from 0 to 133 feet
134 — 266	For line length from 134 to 266 feet
267 — 399	For line length from 267 to 399 feet
400 — 533	For line length from 400 to 533 feet
534 — 655	For line length from 534 to 655 feet

The default for 'long' buildout is 'NotApplicable' while the default for 'short' buildout is '0 — 133'.

channel-group

Syntax [**no**] **channel-group** *channel-group-id*

Context config>port>tdm>ds1>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.

Default None

Description *channel-group-id* — Identifies the channel-group ID number.

Values DS1: 1 — 24
E1: 1 — 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 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 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 config>port>tdm>ds1
config>port>tdm>ds3
config>port>tdm>e1
config>port>tdm>e3

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

TDM DS3/E3	LoopTimed	Default
Channelized ASAP DS3/E3	Yes	node-timed
CES OC-3	Yes	node-timed
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.

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 16 for non-ATM channel groups configured under DS-1, E-1 and for non-ATM E-3 and DS-3 channel/ports.

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** — Use 16 bit checksum for the associated port/channel.

32 — Use 32 bit checksum for the associated port/channel.

down-count

Syntax **down-count** *down-coun*t
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 3

down-count — Specifies the number of keep alive intervals that must pass without receiving a keep alive packet before the link is declared down.

Values 3 — 16

encap-type

Syntax **encap-type** {**atm** | **bcp-null** | **bcp-dot1q** | **ipcp** | **ppp-auto** | **frame-relay** | **wan-mirror** |**cisco-hdlc**}

Context config>port>tdm>ds1>channel-group
config>port>tdm>ds3
config>port>tdm>e1>channel-group
config>port>tdm>e3

Context 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 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 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 will run.

cisco-hdlc — Monitors line status on a serial interface by exchanging keepalive request messages with peer network devices.

cem — On circuit emulation CMAs and MDAs, only the **cem** encap-type is supported. All other values are blocked with an appropriate warning. The **cem** encap-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 DS1: 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 C-HDLC 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 g704

Parameters **g704** — Configure 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 **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-ds1 — 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 E-3 non-ATM: g751 and cannot be changed.
E-3 ATM: 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**}

Context config>port>tdm>ds1>channel-group
config>port>tdm>ds3
config>port>tdm>e1
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 pattern** *pattern*
no idle-payload-fill
idle-payload-fill {**all-ones**}

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 all-ones

Parameters **all-ones —** Defines the 8 bit value to be transmitted as 11111111.

pattern — Transmits a user-defined pattern.

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 all-ones

Parameters **all-ones —** Defines the 8 bit value to be transmitted as 11111111.

pattern — Transmits a user-defined pattern.

insert-single-bit-error

Syntax **insert-single-bit-error**

Context config>port>tdm>ds1
config>port>tdm>e1

Description This command inserts a single bit error for the BERT test.

Default no bit-error-insertion

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

config>port>tdm>e1

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: 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 a internal loopback mode. A 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>e3
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: 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 a 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.

mdl

Syntax **mdl** {**eic** | **lic** | **fic** | **unit** | **pfi** | **port** | **gen**} *mdl-string*
no mdl

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* — specify 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
config>port>tdm>e3

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 —** Specify the MDL path message.

idle-signal — Specify the MDL idle signal message.

test-signal — Specify the MDL test signal message.

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

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 64

Parameters **56 —** Specifies that 56k byte (7-bits per byte) encoding will be used for the associated DS-0 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 — Specify the subrate value for the associated DS-3.

Values 1 — 147 (digital-link)
1 — 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 line & section 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** *timeslots*
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.

Default no timeslots — Non-ATM channel groups.1-24 — Channel groups configured under DS-1 with encap set to ATM.
2-16,18-32 — Channel groups configured under E-1 with encap set to ATM.

Description *timeslots* — 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 — 24 for DS-1 interfaces (the full range is auto-configured for ATM channel groups and cannot be changed).
2 — 32 for E-1 interfaces (the 2 — 16,18 — 32 ranges are auto-configured for ATM channel groups and cannot be changed).

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: 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.

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.

Default No LAGs are defined.

Parameters *lag-id* — The LAG identifier, expressed as a decimal integer.

Values 1 — 800 (7750 SR-c12/4: 1 — 64)

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 XMAs/MDAs. This command applies only to access LAGs.

Default distribute

Parameters Specify the QoS adaptation type:

Values **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 intance 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 IOMXCM as parent of the configured SAP queues and vitual schedulers per LAG member port on that IOMXCM. This internal virtual scheduler limits the total amount of egress bandwidth for all member ports on the IOMXCM 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.

disable-soft-reset-extension

Syntax [**no**] **disable-soft-rest-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.

Parameters **disable-soft-reset-extension —** Disables the automatic extension of BFD timers during an IOM/IMM soft-reset.

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**

per-fp-egr-queuing

Syntax [**no**] **per-fp-egr-queuing**

Context config>lag

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

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.

bfd

Syntax **bfd**

Context config>lag

Description This command creates the bfd context and enables BFD over the associated LAG links.

family

Syntax **family [ipv4 | ipv6]**
no family

Context config>lag>bfd

Description This command is used to specify which address family should be used for the micro-BFD session over the associated LAG links.

Default None

Parameters **ipv4 —** IPv4 encapsulation should be used for the micro-BFD session.

ipv6 — IPv6 encapsulation should 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 *ip-address* — 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:x:d.d.d.d
	x: [0 — FFFF]H
	d: [0 — 255]D

max-admin-down-time

Syntax **max-admin-down-time [*down-interval* | infinite]**
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 no max-admin-down-time

Parameters *down-interval* — Specifies the amount of time, in seconds.

Values -1—3600

infinite — Specifies no end time to forward traffic.

max-setup-time

Syntax **max-setup-time [*up-interval* | infinite]**
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 no max-setup-time

Parameters *up-interval* — Specifies the amount of time, in milliseconds.

Values -1—60000

infinite — Specifies no end time to forward traffic.

multiplier

Syntax **multiplier *multiplier***
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 no multiplier

Parameters *multiplier* — Specifies the multiplier value.

Values 3—20

receive-interval

Syntax **receive-interval *receive-interval***
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 no receive-interval

Parameters *receive-interval* — Specifies the interval value, in milliseconds.

Values 10—100000

Default 100 ms for CPM3 or later, 1 sec for all other

remote-ip-address

Syntax **remote-ip-address *ip-address***
no remote-ip-address

Context config>lag>bfd>family

Description This command is used to specify 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 *ip-address* — 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:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D

transmit-interval

Syntax **transmit-interval *transmit-interval***
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 no transmit-interval

Parameters *transmit-interval* — Specifies the interval value, in milliseconds.

Values 10—100000

Default 100 ms for CPM3 or later, 1 sec for all other

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 no transmit-interval

dynamic-cost

Syntax [**no**] **dynamic-cost**

Context config>lag *lag-id*

Description This command enables OSPF/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 autocost 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 autocost 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 autocost 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 **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 signalling 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 0

Parameters **down** *hold-down-time* **—** Specifies the hold-time for event reporting

Values 0 — 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.

Default no lacp

Parameters Note: 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.

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.

administrative-key *admin-key* **—** 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 — 65535

system-priority *priority* **—** Specifies the system priority.

Values 1 — 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 **coupled**

Parameters **coupled —** TX and RX activate together.

independent — RX activates independent of TX.

lacp-xmit-interval

Syntax **lacp-xmit-interval** {**slow** | **fast**}

Context config>lag

Description This command specifies the interval signaled to the peer and tells the peer at which rate it should transmit.

Default 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** *link-map-profile-id* [**create**]
no link-map-profile *link-map-profile-id*

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.

The **no** form of this command, deletes the specified link map profile.

Default Link-map-profiles are not created by default.

Parameters *link-map-profile-id* — An integer from 1 to 64 that defines a unique lag link map profile on this LAG.

link

Syntax **link** *port-id* {**primary**|**secondary**}
no primary-link

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/network interfaces that use this LAG link map profile.

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.

Default Links are part of a profile.

Notes When a link gets added/deleted, all SAPs/network interfaces that use this link-map-profile may be re-hashed if required.

Parameters *port-id* — A physical port Id in the slot/mda/port format that is an existing member of this LAG.

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.

Options include:

- **discard** – 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** – 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.

The **no** form of this command restores the default failure-mode value.

Default **failure-mode per-link-hash**

port

Syntax **port** *port-id* [*port-id …*] [**priority** *priority*] [**subgroup** *sub-group-id*]
no port *port-id* [*port-id …*]

Context config>lag>port

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 *port-id* — The port ID configured or displayed in the *slot/mda/port* format.

Note that the maximum number of ports in a LAG depends on platform-type, H/W deployed, and SROS S/W 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 1-64 inclusive.

Values slot/mda/port

priority *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 — 65535

subgroup *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 signalled 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 — 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-mda** subgroup is defined based on the MDA. (all ports of the same MDA are assigned to the same subgroup).

port-threshold

Syntax **port-threshold** *value* [**action** {**dynamic-cost** | **down**}
no port-threshold

Context config>lag *lag-id*

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.

The **no** form of this command reverts to the default values.

Default 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 — 63

action {**dynamic-cost** | **down}** — Specifies the action to take if the number of active links in the LAG is at or below the threshold value.

When the **dynamic-cost** action is specified, then 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.

When the **down** action is specified, then 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.

port-type

Syntax **port-type** {**standard** | **hsmda**}

no port-type

Context config>lag

Description This command specifies the type of ports allowed in this LAG.

Parameters **standard —** Allows all non-HSMDA type ports to be added to this LAG.

hsmda — Limits the LAG members to be high-speed MDA (HSMDA) ports only.

port-weight-speed

Syntax **port-weight-speed {1 | 10}**
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:

no port-weight-speed – all LAG links must be of the same speed. Each link weight is 1.

Parameters **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** or **port-weight-speed 10** in service, when all links of the LAG are 1GE or 10GE respectively.

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 or 10GE respectively.

All other configuration changes require shutdown of the LAG and removal of all ports first.

Default no port-weight-speed

selection-criteria

Syntax **selection-criteria** {**highest-count** | **highest-weight** | **best-port**} [**slave-to-partner**] [**subgroup-hold-time** *hold-time*]
no selection-criteria

Context config>lag

Description This command specifies which selection criteria should be used to select the active sub-group.

Default 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).

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.

subgroup-hold-time *hold-time* **—** 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 **not specified** – Equivalent to specifying a value of 0. Specifies no delay and to switchover immediately to a new candidate active sub-group.

Values 0..2000 – Integer specifying the timer value in 10ths of a second.

Values **infinite** – Do not switchover from existing active sub-group if the subgroup remains UP. Manual switchover possible using tools perform lag force command.

standby-signalling

Syntax **standby-signalling** {**lacp** | **power-off**}
no standby-signalling

Context config>lag

Description This command specifies how the state of a member port is signalled to the remote side when the status corresponding to this member port has the **standby** value.

weight-threshold

Syntax **weight-threshold** *value* **action [{dynamic-cost | down}]**
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 no weight-threshold

Parameters *value* — 0..63

action { dynamic-cost | down} — Specifies the action to take if the total weight of active links in the LAG is at or below the threshold value. When the dynamic-cost action is specified then 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. When the down action is specified then the LAG will be brought operationally down if the total weight of operational links is equal to or less than the configured threshold value. The LAG will only be regarded as up once the total weight of operational links exceeds the configured threshold value.

Eth Tunnel Commands

eth-tunnel

Syntax **eth-tunnel** *tunnel-id*
no eth-tunnel

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.

Default no eth-tunnel

Parameters *tunnel-id* — Specifies the tunnel ID.

Values 1 — 64

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** *down-timeout* — Species the eth-tunnel CCM down timers.

Values 0 — 1000 in 100ths of seconds

Default 0

up *up-timeout* — Species the eth-tunnel CCM up timers.

Values 0 — 5000 in 10ths of seconds

Default 20

description

Syntax **description** *long-description-string*
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 *string* — Specifies the text description up to 160 characters in length.

ethernet

Syntax **ethernet**

Context config>eth-tunnel

Description This command is the node where Ethernet parameters can be configured.

encap-type

Syntax **encap-type** {**dot1q**|**qinq**}
no encap-type

Context config>eth-tunnel>ethernet

Description This command configures the encapsulation method.

Parameters **dot1q** — Specifies dot1q encapsulation.

qinq — Specifies qinq encapsulation.

mac

Syntax [**no**] **mac** *ieee-address*

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.

hold-time

Syntax **hold-time**

Context config>eth-tunnel

Description This command configures eth-tunnel dampening timers.

member

Syntax **member down** *time*
no member

Context config>eth-tunnel>hold-time

Description A default MAC address is assigned by the system from the chassis MAC address pool.This command specifies the timer, which controls the delay between detecting that member path is down and reporting it to the G.8031 protection module. If a non-zero value is configured, the CPM will wait for the time specified in the value parameter before reporting it to the G.8031 protection module. Note that this parameter applies only to member path CCM. It does NOT apply to the member port link state. To damp member port link state transitions, use hold-time parameter from the physical member port.

The **no** form of this command sets the hold-time to the default value.

Default no member - the fault will be reported immediately to the protection module.

Parameters *value* — Specifies the hold-time for reporting the failure.

Values 1-1000 centiseconds

lag-emulation

Syntax lag-emulation

Context config>eth-tunnel

Description This command configures eth-tunnel loadsharing parameters.

access

Syntax **access**

Context config>eth-tunnel>lag-emulation

Description This command configures eth-tunnel loadsharing 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** *num-paths*
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

Parameters *num-paths* — Specifies the threshold for the Ethernet Tunnel group.

Values 0 — 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 the default.

Default no path-threshold

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** *time*
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 **n**o form of this command sets the revert-time to the default value.

Default no revert-time – indicates non-revertive behavior.

Parameters *time* — Specifies the re-activation delay in seconds for the primary path.

Values 1 — 720 seconds

path

Syntax [**no**] **path** *path-index*

Context config>eth-tunnel

Description This command configures one of the two paths supported under the Ethernet tunnel. Although the values indicate 1 — 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 need to be un-configured before the path can be deleted.

Default no path

Parameters *path-index* — Specifies the identifier for the path.

Values 1 — 8

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** *vlan-id*
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 specified

Parameters *vlan-id* — specifies the value of the VLAN ID to be used for the control tag.

Values 1 – 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 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 *mep-id* — specifies the maintenance association end point identifier.

Values 1 — 81921

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 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** *priority*
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 *priority* — Specifies the priority of CCM and LTM messages.

Values 0 — 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 *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

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 remErrXcon

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
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** *mac-address*
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 *mac-address — 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.

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 **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**vlan** *vlan-id*]
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 *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 81921

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

vlan-id — 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 — 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** [[*level* [*level* ...]]
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 *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

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

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 *priority-value* — Specify the priority value of the AIS messages originated by the node.

Values 0 — 7

Default 7

ccm-enable

Syntax **[no] ccm-enable**

Context config>port>ethernet>eth-cfm>mep
config>lag>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** *priority*
no ccm-ltm-priority

Context config>port>ethernet>eth-cfm>mep>
config>lag>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 *priority* — Specifies the priority value

Values 0 — 7

Default 7

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*
no ccm-padding-size

Context config>eth-tunnel>path>eth-cfm>mep

Description This command inserts additional padding in the CCM packets.

The **no** form of the command reverts to the default.

Parameters *ccm-padding* — Specifies the additional padding in the CCM packets.

Values 3 — 1500 octets

ccm-tlv-ignore

Syntax **ccm-tlv-ignore** [**port-status**] [**interface-status**]
no ccm-tlv-ignore

Context config>port>ethernet>eth-cfm>mep
config>lag>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 —** Ignore the port status TLV on reception.

interface-status — ignore the interface status TLV on reception.

collect-lmm-stats

Syntax **collect-lmm-stats**
[no] collect-lmm-stats

Context config>port>ethernet>eth-cfm>mep
config>router>if>eth-cfm>mep
config>lag>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 lmm 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-lmm-stats

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>port>ethernet>eth-cfm>mep
config>lag>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 *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

The **no** form of the command disables eth-test capabilities.

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>eth-ring>path>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default 1

Parameters *bit-errors* — Specifies the lowest priority defect.

Values 0 — 11840

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>port>ethernet>eth-cfm>mep>eth-test
config>lag>eth-cfm>mep>eth-test
config>router>if>eth-cfm>mep>eth-test

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.

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>port>ethernet>eth-cfm>mep>eth-test
config>lag>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, well enabled to do so, causes a network reaction.

Default macRemErrXcon

Values

allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
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** *mac-address*
no mac-address

Context config>port>ethernet>eth-cfm>mep
config>lag>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 *mac-address* — 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** *seconds*

Context config>eth-tunnel>path>eth-cfm>mep

Description This command enables one way delay threshold time limit.

Default 3 seconds

Parameters *seconds* — Specifies the value, in seconds, for the threshold.

Values 0 — 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.

Parameters *accept* — Share fate with the facility tunnel MEP

ignore — Do not share fate with the facility tunnel MEP

Default **ignore** (Service Level)

accept (SAP Level for Epipe and VPLS)

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.

NOTE: 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.

Values <*boot-env*|*config*> : keywords

Refer to the 7750 SR OS Basic System Configuration Guide.

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 enabled

Parameters **boot-env —** Synchronizes all files required for the boot process (loader, BOF, images, and configuration files.

config — Synchronize only the primary, secondary, and tertiary configuration files.

Default config

bgp-multi-homing

Syntax **bgp-multi-homing**

Context config>redundancy

Description This command configures BGP multi-homing parameters.

boot-timer

Syntax **boot-timer** *seconds*
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 10 sec

Parameters *seconds* — Specifies the timer, in seconds.

Values 1..100

site-activation-timer

Syntax **site-activation-timer** *seconds*
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 2 seconds

Parameters *seconds* — Specifies the timer, in seconds.

Values 1..100

site-min-down-timer

Syntax **site-min-down-timer** *min-down-time*
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 default value.

Default no site-min-down-timer

Parameters *min-down-time* — Specifies the time, in seconds, that a BGP multi-homing site remains operationally down after a transition from up to down.

Values 1— 100 seconds

Default 0 seconds

multi-chassis

Syntax **multi-chassis**

Context config>redundancy

Description This command enables the context to configure multi-chassis parameters.

peer

Syntax [**no**] **peer** *ip-address* **create**

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 *ip-address* — 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:x:d.d.d.d
	x: [0 — FFFF]H
	d: [0 — 255]D

create — Mandatory keyword specifies to create the peer.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**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.

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, etc.), 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 55 (hash2-key) characters in length (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 a non-encrypted, 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 then the key value alone, this means that hash2 encrypted variable cannot be copied and pasted. If the hash or

hash2 parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

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.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>redundancy>multi-chassis>peer>mc-ep
config>router>rsvp
config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor
config>redundancy>multi-chassis>peer>mc-ep

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 300

Parameters *interval* — Specifies the boot timer interval.

Values 1 — 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 3

Parameters *multiplier* — Specifies the hold time applied on neighbor failure.

Values 2 — 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 5 (0.5s)

Parameters *interval* — The time interval expressed in deci-seconds.

Values 5 — 500 (tenths of a second)

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 highest value is chosen to be the Master. If system-priority are equal 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 default

Default 0

Parameters *value* — Specifies the priority assigned to the local MC-EP peer.

Values 1— 255

MC LAG Commands

mc-lag

Syntax [**no**] **mc-lag**

Context config>redundancy>multi-chassis>peer>mc-lag

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.

hold-on-neighbor-failure

Syntax **hold-on-neighbor-failure** *multiplier*
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 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 — 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 1s (10 hundreds of milliseconds means interval value of 10)

Parameters *interval* — The time interval expressed in deci-seconds

Values 5 — 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*.

Default none

Parameters *lag-id* — The LAG identifier, expressed as a decimal 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 — 800 (7750 SR-c12/4: 1 — 64)

lacp-key ***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 — 65535

system-id ***system-id*** **—** Specifies a 6 byte value expressed in the same notation as MAC address

Values xx:xx:xx:xx:xx:xx - xx [00..FF]

remote-lag ***lag-id*** **—** Specifies the LAG ID on the remote system.

Values 1 — 800

system-priority *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 — 65535

source-bmac-lsb *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 — 65535 or xx-xx or xx:xx

source-address

Syntax **source-address** *ip-address*
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 *ip-address* — Specifies the source address used to communicate with the multi-chassis peer.

sync

Syntax [**no**] **sync**

Context config>redundancy>multi-chassis>peer

Description This command enables the context to configure synchronization parameters.

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 specifies whether MLD snooping information should be synchronized with the multi-chassis peer.

Default no mld-snooping

port

Syntax **port** [*port-id* | *lag-id*] [**sync-tag** *sync-tag*]
no port [*port-id* | *lag-id*]

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 *port-id* — Specifies the port to be synchronized with the multi-chassis peer.

lag-id — Specifies the LAG ID to be synchronized with the multi-chassis peer.

sync-tag *sync-tag* — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

range

Syntax **range** *encap-range* **sync-tag** *sync-tag*
no range *encap-range*

Context config>redundancy>multi-chassis>peer>sync>port

Description This command configures a range of encapsulation values.

Parameters **Values** encap-range

Specifies a range of encapsulation values on a port to be synchronized with a multi-chassis peer.

Values		
	Dot1Q	*start-vlan-end-vlan*
	QinQ	Q1.*start-vlan*-Q1.*end-vlan*

sync-tag *sync-tag* — Specifies a synchronization tag up to 32 characters in length 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 [**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

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.

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 **Values** sync-tag

Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

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.

Parameters *ip-address* — Specifies the destination IP address.

interface

Syntax **interface** *ip-int-name*
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.

service-id

Syntax **service-id** *service-id*
no service-id

Context config>redundancy>mc>peer>mcr>ring>ibc
config>redundancy>mc>peer>mcr>node>cv

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.

The **no** form of the command removes the service-id from the IBC configuration.

Parameters *service-id* — Specifies the service ID if the interface.

Values *service-id*: 1 — 2147483647

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** *vlan-range*

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 *vlan-range* — Specifies the VLAN range.

Values [0 — 4094] — [0 — 4094]

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** *ring-node-name* [**create**]
no ring-node *ring-node-name*

Context config>redundancy>mc>peer>mcr>ring

Description This command specifies the unique name of a multi-chassis ring access node.

Parameters *ring-node-name* — Specifies the unique name of a multi-chassis ring access node.

create — Keyword used to create the ring node instance. The **create** keyword requirement can be enabled/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** *interval*
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 5

Parameters *interval* — Specifies the polling interval, in minutes.

Values 1 — 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.

Default no service-id

Parameters *service-id* — Specifies the service ID of the SAP.

Values 1 — 2147483647

Values *service-id*: 1 — 2147483647

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 *ip-address* — Specifies the source IP address.

src-mac

Syntax **src-mac** *ieee-address*
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 zeroes (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 *ieee-address* — Specifies the source MAC address.

vlan

Syntax **vlan** [*vlan-encap*]
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 *vlan-encap* — Specifies the VLAN tag.

Values

vlan-encap:	dot1q	qtag
	qinq	qtag1.qtag2
	qtag	0 — 4094
	qtag1	1 — 4094
	qtag2	0 — 4094

Forwarding Plane Commands

fp

Syntax **fp** [*fp-number*]

Context config>card

Description This command enables the context to configure multicast path management commands for IOM-3 ingress multicast management. Ingress multicast management manages multicast switch fabric paths which are forwarding plane specific. On IOM-1 and IOM-2, each MDA has a dedicated forwarding plane and so have dedicated multicast paths to the switch fabric allowing the multicast management to be defined per MDA. IOM-3 has a single forwarding plane shared by two MDAs. The fp node simplifies ingress multicast management on IOM-3.

While IOM-3 only has a single forwarding plane. In future releases, to accommodate multiple forwarding planes, each forwarding plane will be assigned a value. 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.

Parameters *fp-number* — The fp-number parameter is optional following the **fp** command. If omitted, the system assumes forwarding plane number 1.

Values 1

Default 1

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>card>fp

Description This command specifies the protocol name to be monitored by Distributed CPU Protection Policy.

egress

Syntax **egress**

Context config>card>fp

Description This command enables the egress **fp** node that contains the multicast path management configuration commands for IOM-3 ingress multicast management.

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 IOM3-XP forwarding plane.

buffer-allocation

Syntax **buffer-allocation min** *percentage* **max** *percentage*
no buffer-allocation

Context config>card>fp>egress>max-wred-control

Description The buffer-allocation command defines the amount of IOM3-XP buffers that will be set aside for WRED queue buffer pools. **Note** that the **min** *percentage* and max *percentage* parameters must be set to the same value. The IOM3-XP 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 and the size if further fine-tuned by the forwarding class oversubscription factors.

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. Further oversubscription controls are described within the resv-cbs command later in this document.

The WRED mega-pool is allowed to expand between the min and max percent of total forwarding plane buffers based on the sum of the WRED queue sizes and the WRED oversubscription factors. As the WRED mega-pool grows, the number of buffers available to the default mega-pool will shrink. If the WRED mega-pool shrinks, the default mega-pool will grow accordingly. When min and max are defined as the same value, the WRED mega-pool size will not fluctuate and the oversubscription factors will have no effect.

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.

FC MBS Oversubscription Factors and WRED Mega-Pool Sizing

Each WRED queue in a SAP egress QoS policy is created on an egress IOM3-XP when the policy is applied to an egress SAP on the IOM and at least one forwarding class is mapped to the queue. For WRED queue buffer management purposes, each forwarding class is configured with an MBS oversubscription factor (OSF) on the IOM using the **osf** command. The MBS oversubscription factor is used by the system as a pro-

visioning parameter that defines the acceptable level of oversubscription between the sum of the maximum buffer sizes (mbs) of the WRED queues for a given class and the number of buffers for that class in the WRED mega-pool. Since multiple forwarding classes may be mapped to the same queue, the oversubscription factor associated with the highest forwarding class mapped is used for dynamically sizing the WRED mega-pool.

As an example, when a WRED queue is configured with the following attributes:

- MBS equal to 10Kbytes
- AF as the highest forwarding class mapped

And the forwarding plane on the IOM3-XP is configured with the following WRED limits:

- Current WRED mega-pool is sized at 500Kbytes
- AF MBS oversubscription factor is 2 (2:1)

The system will increase the WRED mega-pool size to 505Kbytes (increase of 10Kbytes/2) as long as the maximum buffer allocation percentage equates to a value equal to or greater than 505Kbytes. (If not, the WRED mega-pool will be capped at the maximum level.)

The **no** form of the command immediately restores the default min and max percentage values for sizing the WRED mega-pool.

Parameters **min** *percent-of-total* **—** This required keyword defines the minimum percentage of total IOM3-XP queue buffers that will be applied to the WRED mega-pool. The value given for percent-of-total must be less than or equal to the value given for the **max** *percent-of-total*. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%).

Values 0.00 — 99.99

Default 25.00

max *percent-of-total* **—** This required keyword defines the maximum percentage of total IOM3-XP queue buffers that may be applied to the WRED mega-pool. The value given for percent-of-total must be greater than or equal to the value given for the **min** *percent-of-total*. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%).

Values 0.01 — 99.99

Default 25.00

resv-cbs

Syntax **resv-cbs min** *percentage* **max** *percentage*
no resv-cbs

Context config>card>fp>egress>max-wred-control

Description This command defines the amount of IOM3-XP buffers within the WRED mega-pool that will be set aside for WRED queues operating within their configured CBS thresholds. **Note** that the **min** *percentage* and **max** *percentage* parameters must be set to the same value. The IOM3-XP 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. Further control over shared congestion is defined later in this document under the slope-policy command.

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. Forwarding class cbs-factor settings are used in the same way as the mbs-factor parameters to move the actual reserved size between the minimum and maximum thresholds according to appropriate oversubscription factors that are applied to the sum of the WRED queue CBS values.

When min and max are defined as the same value, the WRED mega-pool size will not fluctuate and the oversubscription factors will have no effect.

FC CBS Oversubscription Factors and WRED CBS Reserve Sizing

Each WRED queue in a SAP egress QoS policy is created on an egress IOM3-XP when the policy is applied to an egress SAP on the IOM and at least one forwarding class is mapped to the queue. For WRED queue CBS buffer management purposes, each forwarding class is configured with a CBS oversubscription factor (OSF) on the IOM using the **osf** command. The CBS oversubscription factor is used by the system as a provisioning parameter that defines the acceptable level of oversubscription between the sum of the committed buffer sizes (CBS) of the WRED queues for a given class and the number of buffers for that class that should be placed in the WRED mega-pool CBS reserve. Since multiple forwarding classes may be mapped to the same queue, the oversubscription factor associated with the highest forwarding class mapped is used for dynamically sizing the WRED mega-pool CBS reserve.

As an example, when a WRED queue is configured with the following attributes:

- CBS equal to 6Kbytes
- AF as the highest forwarding class mapped

And the forwarding plane on theIOM3-XP is configured with the following WRED limits:

- Current WRED mega-pool CBS reserve is sized at 100Kbytes
- AF CBS oversubscription factor is 2 (2:1)

The system will increase the WRED mega-pool CBS reserve size to 103Kbytes (increase of 6Kbytes/2) as long as the maximum buffer allocation percentage for resv-cbs equates to a value equal to or greater than 103Kbytes. (If not, the WRED mega-pool CBS reserve will be capped at the maximum level.)

The **no** form of the command immediately restores the default min and max percentage values for sizing the WRED mega-pool CBS reserve.

Parameters **min** *percent-of-total* — This required keyword defines the minimum percentage of the IOM3-XP WRED mega-pool buffers that will be applied to the CBS reserve. The value given for percent-of-wred must be less than or equal to the value given for the max percent-of-wred. Percentages are defined with an accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%).

Values 0.00 — 99.99

Default 25.00

max *percent-of-total* — This 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 percent-of-wred must be greater than or equal to the value given for the min percent-of-wred. Percentages are defined with an

accuracy of hundredths of a percent in the nn.nn format (15.65 = 15.65%).

Values 0.01 — 99.99

Default 25.00

slope-policy

Syntax **slope-policy** *slope-policy-name*
no slope-policy

Context config>card>fp>egress>max-wred-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 in-profile, the packet will be associated with the high priority slope. Out-of-profile packets are associated with the low priority 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.

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.

Parameters *slope-policy-name* — This required parameter 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.

Default When not defined, the default slope policy is used

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 - mdaHiBwMulticastAlarm). 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.

This feature is supported on the 7750 SR-7 and 7750 SR-12.

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.

group *group-id* **—** Specifies the logical MSFP group for the MDA. MDAs configured with the same *group-id* will be placed on the same MSFP.

Values 0 — 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.

shutdown

Syntax [**no**] **shutdown**

Context config>card>fp>egress>max-wred-control

Description This command enables or disables egress WRED queue support on the IOM. By default, WRED queue support is disabled (shutdown). While disabled, the various wred-queue-control commands may be executed on the IOM and SAP egress QoS policies with wred-queue enabled may be applied to egress SAPs. The IOM 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 IOM3-XP.

ingress

Syntax **ingress**

Context config>card>fp

Description The ingress CLI node within the **fp** node contains the multicast path management configuration commands for IOM-3 ingress multicast management. The **bandwidth-policy** command is supported within the ingress node.

stable-pool-sizing

Syntax [**no**] **stable-pool-sizing**

Context config>card>fp

Description The stable-pool-sizing command is used to provide 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.

The **no** stable-pool-sizing command is used to disable stable pool sizing on a forwarding plane. Existing buffer pools will be resized according to normal pool sizing behavior.

access

Syntax **access**

Context config>card>fp>ingress

Description This CLI node contains the access forwarding-plane parameters.

queue-group

Syntax **queue-group** *queue-group-name* **instance** *instance-id* [**create**]
no queue-group

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 queue-group-name and **instance** *instance-id* 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.

Default none

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 in length. 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 — 16383

create — Keyword used to associate the queue group. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

queue-group

Syntax **queue-group** *queue-group-name* **instance** *instance-id*
no queue-group

Context config>card>fp>ingress>network

Description This command is used to create 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.

Default none

Parameters *queue-group-name* — Specifies the name of the queue group template up to 32 characters in length.

instance-id — pecifies the identification of a specific instance of the queue-group.

Values 1— 16384

accounting-policy

Syntax **accounting-policy** *policy-name*
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 *policy-name* — Specifies the name of the accounting policy to use for the queue-group.

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-policy

Syntax **policer-control-policy** *policy-name*
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 policies are specified by default. You must explicitly specify a policy.

Parameters *policy-name* — Specifies the name of the policer-control policy to use for the queue-group.

ingress-buffer-allocation

Syntax **ingress-buffer-allocation** *hundredths-of-a-percent*
no ingress-buffer-allocation

Context config>card>fp>ingress

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.

NOTE: This feature is supported on all 50G FP2-based line cards and 100G/200G FP3-based line cards.

The **no** form of this command returns the ingress buffer allocation to the default value.

Default The default value is 50.00%, which emulates the legacy behavior.

max-rate

Syntax **max-rate** {*kilobits-per-second* | **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.

Default max

Parameters *kilobits-per-second* — Defining a kilobits-per-second value is mutually exclusive with the max parameter. 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 Integer 0 – 2000000000

max — The **max** parameter 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.

no max-rate — The **no max-rate** command returns the policer-control-policy's parent policer maximum rate to max.

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.

Default None.

min-thresh-separation

Syntax **min-thresh-separation** *size* [**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* [**bytes** | **kilobytes**] — 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 – 16777216

Default none

[**bytes** | **kilobytes**] — 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 **priority** *level*

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 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.

Default None.

mbs-contribution

Syntax **mbs-contribution** *size* [**bytes | kilobytes**] [**fixed**]
no mbs-contribution

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 The **mbs-contribution** command is used to configure 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 Mbps (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 Mbps.
- 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 Mbps, but only 8 Mbps is available,
- Based on weights, the children's FIR rates are set as follows:

	FIR Rate	FIR MBS
Child 1	4 Mbps	10 Kbytes
Child 2	3 Mbps	10 Kbytes
Child 3	1 Mbps	10 Kbytes

The 12 Mbps of the higher priority traffic and the 8 Mbps of fair traffic equal the 20 Mbps decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mbps of the parent policer's decrement rate, leaving 8 Mbps of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 Kbytes above 4 Mbps,
- The burst tolerance of child 2 is based on 10 Kbytes above 3 Mbps,
- The burst tolerance of child 3 is based on 10 Kbytes above 1 Mbps.

If all three children burst simultaneously (unlikely), they will consume 30 Kbytes above 8 Mbps. 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 Policer 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.

Parameters *size* [**bytes** | **kilobytes**] — 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 — 16777216

Default none

bytes | **kilobytes**: — 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**

fixed — The optional fixed keyword is used to force the inclusion of the defined **mbs-contribution** value (or an override value defined on the SAP or sla-profile) in the parent policer's discard threshold calculations. If the **mbs-contribution** command is executed without the **fixed** keyword, the fixed calculation behavior for the priority level is removed.

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.

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 is used to remove any existing policer overrides.

Default no policer-overrides

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 is used in the sap-ingress and sap-egress QoS policies to create, modify or delete 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 will fail.

The **no** form of this command is used to delete 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* — 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—32

stat-mode

Syntax **stat-mode** {**no-stats** | **minimal** | **offered-profile-no-cir** | **offered-priority-no-cir** | **offered-limited-profile-cir** | **offered-profile-cir** | **offered-priority-cir** | **offered-total-cir**}
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 is used to configure 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 sys-**

tem-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 **no-stats —** Counter resource allocation:0

The policer does not have any forwarding plane counters allocated and cannot provide offered, discard and forward statistics. A policer using no-stats cannot be a child to a parent policer and the policer's parent command will fail.

When **collect-stats** is enabled, the lack of counters causes the system to generate the following statistics:

a. offered-in = 0

b. offered-out = 0

c. discard-in = 0

d. discard-out = 0

e. forward-in = 0

f. forward-out= 0

Counter 0 indicates that the accounting statistic returns a value of zero.

minimal — Counter resource allocation:1

The default **stat-mode** for a policer is **minimal**. The **minimal** mode allocates 1 forwarding plane offered counter and one traffic manager discard counter. The forwarding counter is derived by subtracting the discard counter from the offered counter. The counters do not differentiate possible offered types (profile or priority) and do not count green or yellow output. This does not prevent the policer from supporting different offered packet types and does not prevent the policer from supporting a CIR rate.

This counter mode is useful when only the most basic accounting information is required.

The counters are used in the following manner:

1. 'offered = profile in/out, priority high/low

2. 'discarded = Same as 1

3. 'forwarded= Derived from 1 - 2

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out= 0

c. discard-in = 2

d. discard-out= 0

e. forward-in = 3

f. 'orward-out= 0

Counter 0 indicates that the accounting statistic returns a value of zero.

With **minimal** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 1

ii. offered-out= 0

iii. offered-undefined= 0

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-no-cir — Counter resource allocation:2

The **offered-profile-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-profile-no-cir** mode is most useful when the policer is receiving only in-profile and out-of-profile pre-marked (and trusted) packets. It is expected that in this instance a CIR rate will not be defined since all packet are already pre-marked. This mode does not prevent the policer from receiving un-trusted (color undefined) nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-in = profile in

2. offered-out= profile out, priority high/low

3. dropped-in= Same as 1

4. dropped-out= Same as 2

5. forwarded-in= Derived from 1 - 3

6. forwarded-out= Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out= 2

c. discard-in = 3

d. discard-out= 4

e. forward-in = 5

f. forward-out= 6

With **offered-profile-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 1

ii. offered-out= 2

iii. offered-undefined= 0

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-no-cir — Counter resource allocation:2

The **offered-priority-no-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-priority-no-cir** mode is most useful when the policer is receiving only un-trusted packets and the ingress priority high and priority low classification options are being used without a CIR profiling rate defined. This mode does not prevent the policer from receiving trusted packets that are pre-marked in-profile or out-of-profile nor does it prevent the policer from being configured with a CIR rate.

The counters are used in the following manner:

1. offered-high = profile in, priority high

2. offered-low= profile out, priority low

3. dropped-high= Same as 1

4. dropped-low= Same as 2

5. forwarded-high= Derived from 1 - 3

6. forwarded-low= Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-high= 1

b. offered-low= 2

c. discard-high= 3

d. discard-low= 4

e. forward-high= 5

f. forward-low= 6

With **offered-priority-no-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1

ii. offered-low= 2

iii. offered-undefined= 0

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-limited-profile-cir — Counter resource allocation:3

The **offered-limitied-profile-cir** mode allocates three forwarding plane offered counters and three traffic manager discard counters.

The **offered-limited-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile (profile out but no profile in) traffic and un-trusted packets are being applied to a defined CIR profiling rate. This mode does not prevent the policer from receiving trusted in-profile packets.

The counters are used in the following manner:

1. offered-undefined-that-turned-green= profile in, priority high/low

2. offered-undefined-that-turned-yellow-or-red= priority high/low

3. offered-out-that-stayed-yellow-or-turned-red= profile out

4. dropped-undefined-that-turned-green= Same as 1

5. dropped-undefined-that-turned-yellow-or-red= Same as 2

6. dropped-out-that-turned-yellow-or-red= Same as 3

7. forwarded-undefined-that-turned-green= Derived from 1 - 4

8. forwarded-undefined-that-turned-yellow= Derived from 2 - 5

9. forwarded-out-that-turned-yellow= Derived from 3 - 6

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 0

b. offered-out= 1 + 2 + 3

c. discard-in = 0

d. discard-out= 4 + 5 + 6

e. forward-in = 7

f. 'orward-out= 8 + 9

With **offered-limited-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-in = 0

ii.'offered-out= 3

iii.'offered-undefined= 1 + 2

iv. offered-managed= 0(IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-profile-cir — Counter resource allocation:4

The **offered-profile-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-profile-cir** mode is most useful when the policer is receiving trusted out-of-profile and in-profile traffic and is also receiving un-trusted packets that are being applied to a defined CIR profiling rate. This mode differs from **offered-limited-profile-cir** mode in that it expects both trusted in-profile and out-of-profile packets while still performing CIR profiling on packets with un-trusted markings. It

is expected that in most cases where both trusted and un-trusted packets are received, the predominate case will not include trusted in-profile packets making the offered-limited-profile-cir accounting mode acceptable.

The counters are used in the following manner:

1. offered-in-that-stayed-green-or-turned-red= profile in

2. offered-undefined-that-turned-green= priority high/low

3. offered-undefined-that-turned-yellow-or-red= priority high/low

4. offered-out-that-stayed-yellow-or-turned-red= profile out

5. dropped-in-that-stayed-green-or-turned-red= Same as 1

6. dropped-undefined-that-turned-green= Same as 2

7. dropped-undefined-that-turned-yellow-or-red= Same as 3

8. dropped-out-that-turned-yellow-or-red= Same as 4

9. forwarded-in-that-stayed-green= Derived from 1 - 5

10. forwarded-undefined-that-turned-green= Derived from 2 - 6

11. forwarded-undefined-that-turned-yellow= Derived from 3 - 7

12. forwarded-out-that-turned-yellow= Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1

b. offered-out= 2 + 3 + 4

c. discard-in = 5 + 6

d. discard-out= 7 + 8

e. forward-in = 9 + 10

f. forward-out= 11 + 12

With **offered-profile-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1

ii. offered-low= 4

iii. offered-undefined= 2 + 3

iv. offered-managed= 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-priority-cir — Counter resource allocation:4

The **offered-priority-cir** mode allocates four forwarding plane offered counters and four traffic manager discard counters.

The **offered-priority-cir** mode is most useful when the policer is receiving only un-trusted packets that are being classified as high priority or low priority and are being applied to a defined CIR profiling rate.

This mode differs from **offered-profile-cir** mode in that it does not expect trusted in-profile and out-of-profile packets but does not exclude the ability of the policer to receive them.

The counters are used in the following manner:

1. offered-high-that-turned-green= profile in, priority high

2. offered-high-that-turned-yellow-or-red= profile in, priority high

3. offered-low-that-turned-green= profile out, priority low

4. offered-low-that-turned-yellow-or-red= profile out, priority low

5. dropped-high-that-turned-green= Same as 1

6. dropped-high-that-turned-yellow-or-red= Same as 2

7. dropped-low-that-turned-green= Same as 3

8. dropped-low-that-turned-yellow-or-red= Same as 4

9. forwarded-high-that-turned-green= Derived from 1 - 5

10. forwarded-high-that-turned-yellow= Derived from 2 - 6

11. forwarded-low-that-turned-green= Derived from 3 - 7

12. forwarded-low-that-turned-yellow= Derived from 4 - 8

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-high= 1 + 2

b. offered-low= 3 + 4

c. discard-in = 5 + 7

d. discard-out= 6 + 8

e. forward-in = 9 + 11

f. forward-out= 10 + 12

With **offered-priority-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1 + 2

ii. offered-low= 3 + 4

iii. offered-undefined= 0

iv. offered-managed= 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

offered-total-cir — Counter resource allocation:2

The **offered-total-cir** mode allocates two forwarding plane offered counters and two traffic manager discard counters.

The **offered-total-cir** mode is most useful when the policer is not receiving trusted in-profile or out-of-profile traffic and both high and low priority classifications are not being used on the un-trusted packets and the offered packets are being applied to a defined CIR profiling rate. This mode does not prevent

the policer from receiving trusted in-profile or out-of-profile packets and does not prevent the use of priority high or low classifications on the un-trusted packets.

The counters are used in the following manner:

1. offered-that-turned-green= profile in/out, priority high/low

2. offered- that-turned-yellow-or-red= profile in/out, priority high/low

3. dropped-offered-that-turned-green= Same as 1

4. dropped-offered-that-turned-yellow-or-red= Same as 2

5. forwarded-offered-that-turned-green= Derived from 1 - 3

6. forwarded-offered-that-turned-yellow= Derived from 2 - 4

When **collect-stats** is enabled, the counters are used by the system to generate the following statistics:

a. offered-in = 1 + 2

b. offered-out= 0

c. discard-in = 3

d. discard-out= 4

e. forward-in = 5

f. forward-out= 6

Counter 0 indicates that the accounting statistic returns a value of zero.

With **offered-total-cir** enabled as the policer **stat-mode**, the SAP offered stats for the policer returned via MIB query and CLI show commands will return the following values:

i. offered-high= 1 + 2

ii. offered-low= 0

iii. offered-undefined= 0

iv. offered-managed= 0 (IMPM managed packets are not redirected from the policer)

Counter 0 indicates that the SAP policer statistic returns a value of zero.

rate

Syntax **rate** {**max** | **kilobits-per-second**} [**cir** {**max** | **kilobits-per-second**}]
no rate

Context config>card>fp>ingress>access>qgrp>policer-over>plcr
config>card>fp>ingress>network>qgrp>policer-over>plcr

Description This command is used to configure 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 it's 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 Kbps (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 is used to restore the default metering and profiling rate to a policer.

Parameters {**max** | *kilobits-per-second*} — Specifying the keyword **max** or an explicit *kilobits-per-second* 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—2000000000

cir {**max** | *kilobits-per-second*} — The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit *kilobits-per-second* 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 Kbps. 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 CPIR used is equivalent to max.

Values **max** or 0—2000000000

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 is used to configure 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.

Default **none**

Parameters *size* [**bytes** | **kilobytes**] — 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 **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 — 16777216

Default **kilobyte**

mbs

Syntax **mbs** {*size* [**bytes** | **kilobytes**] | **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 is used to configure 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.

Default None

Parameters *size* [**bytes** | **kilobytes**] — The *size* 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.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 — 16777216

Default **kilobyte**

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
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 is used to modify 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 the 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 is used to remove 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 — 31

Default None

subtract *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—64

Default None

mcast-path-management

Syntax **mcast-path-management**

Context config>card>fp>ingress
config>card>mda>ingress

Description This CLI node contains the forwarding plane or MDA settings for ingress multicast path management. Enter the node to configure the bandwidth-policy, the individual path bandwidth overrides and the administrative state of ingress multicast path management.

bandwidth-policy

Syntax **bandwidth-policy** *policy-name*
no bandwidth-policy

Context config>card>fp>ingress>mcast-path-management
config>card>mda>ingress>mcast-path-management

Description This command is used to explicitly associate 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 *policy-name* — 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.

Values Any existing bandwidth policy name

Default default

primary-override

Syntax **primary-override**

Context config>card>mda>ingress>mcast-mgmt

Description This command enables the context to configure MDA ingress multicast path-limit overrides.

The path override CLI nodes are not supported on IOM-3.

secondary-override

Syntax **secondary-override**

Context config>card>mda>ingress>mcast-mgmt

Description This command enables the context to configure MDA ingress multicast path-limit overrides.

The path override CLI nodes are not supported on IOM-3.

ancillary-override

Syntax **ancillary-override**

Context config>card>mda>ingress>mcast-mgmt

Description This command enables the context to configure MDA ingress multicast path-limit overrides.

path-limit

Syntax **path-limit** *megabits-per-second*
no path-limit

Context config>card>mda>ingress>mcast-mgmt>primary-override
config>card>mda>ingress>mcast-mgmt>secondary-override
config>card>mda>ingress>mcast-mgmt>ancillary-override

Description The path-limit command is used to override the path limits contained in the bandwidth policy associated with the MDA. The path limits are used to give the upper limit that multicast channels may use on each path.

The path-limit commands are not supported on IOM-3.

The no form of the command removes a path limit override from an ingress multicast path and restore the path limit defined in the bandwidth policy associated with the MDA.

Parameters *megabits-per-second* — The megabits-per-second parameter is required when executing the path-limit com-

mand and is expressed as an integer representing multiples of 1,000,000 bits per second.

Values Primary-override: 1 to 2000
Secondary-override: 1 to 2000
Ancillary-override: 1 to 5000

Default None

cpm

Syntax **cpm**

Context tools>dump>mcast-path-mgr

Description This command dumps multicast path manager CPM information.

Sample Output

```
*A:Dut-C# tools dump mcast-path-mgr cpm
McPathMgr[10][0]: 0x763a52c0 blkHoleEval 0
     pPath     swPlaneID     pathType      availBw    pathLimit
inUseBw    maxUsedBw numSGs
0x763a54c8              2    secondary      1800000
1800000            0            0      0
0x763a56c0              1      primary      1039959      2000000
960041       960041      6
0x763a58b8             15      primary       879910      2000000
1120090      1120090      7
0x763a5ab0             14      primary       879908      2000000
1120092      1120092      7
0x763a5ca8             13      primary       880007      2000000
1119993      1119993      7
0x763a5ea0             12      primary       880172      2000000
...
0x763a7448              0         none            0
0            0            0      0
0x763a7640              0    blackhole            0
0            0            0      0
McPathMgr[8][0]: 0x7639a9d8 blkHoleEval 0
     pPath     swPlaneID     pathType      availBw    pathLimit
inUseBw    maxUsedBw numSGs
0x7639abe0              1    secondary      1800000
1800000            0            0      0
0x7639add8             15      primary      2000000
2000000            0            0      0
0x7639afd0             14      primary      2000000
...0x7639cd58             0    blackhole           0
0            0            0      0
McPathMgr[9][0]: 0x76398420 blkHoleEval 0
     pPath     swPlaneID     pathType      availBw    pathLimit
inUseBw    maxUsedBw numSGs
0x76398628             15    secondary      1800000
1800000            0            0      0
0x76398820             14      primary      2000000
2000000            0            0      0
```

```
0x76398a18            13      primary      2000000
2000000            0            0      0
...
0x7639a7a0             0    blackhole            0
0           0            0      0
SwPlane[0]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98ba320      2000000      2000000            0      2000000
1800000            0      1800000
SwPlane[1]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98ba390      2000000      2000000       960041      1039959
1800000            0      1039959
###################################

stype inst             src             grp currBw pathBw pref repl path exp
    0    1      10.10.6.33      227.0.0.23 159891 159891    0    0    P   N
    0    1      10.10.4.10       225.0.0.0 159990 159990    0    0    P   N
    0    1      10.10.4.27      225.0.0.17 159990 159990    0    0    P   N
    0    1      10.10.4.43      225.0.0.33 159993 159993    0    0    P   N
    0    1      10.10.6.47      227.0.0.37 160049 160049    0    0    P   N
    0    1      10.10.4.59      225.0.0.49 160128 160128    0    0    P   N
SwPlane[2]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98ba400      2000000      2000000      1119789       880211
1800000            0       880211
###################################
...
###################################

stype inst             src             grp currBw pathBw pref repl path exp
    0    1      10.10.6.29      227.0.0.19 159891 159891    0    0    P   N
    0    1      10.10.4.28      225.0.0.18 159989 159989    0    0    P   N
    0    1      10.10.4.11       225.0.0.1 159990 159990    0    0    P   N
    0    1      10.10.4.41      225.0.0.31 159992 159992    0    0    P   N
    0    1      10.10.6.43      227.0.0.33 160049 160049    0    0    P   N
    0    1      10.10.6.58      227.0.0.48 160052 160052    0    0    P   N
    0    1      10.10.4.55      225.0.0.45 160127 160127    0    0    P   N
SwPlane[16]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98baa20      2000000      2000000            0      2000000
1800000            0      1800000
SwPlane[17]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98baa90      2000000      2000000            0      2000000
1800000            0      1800000
SwPlane[18]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98bab00      2000000      2000000            0      2000000
1800000            0      1800000
SwPlane[19]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
```

```
0x98bab70      2000000      2000000            0      2000000
1800000            0      1800000
SwPlane[20]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
0x98babe0      2000000      2000000            0      2000000
1800000            0      1800000
SwPlane[21]
  pSwPlane      totalBw        priBw   priInUseBw   priAvailBw
secBw   secInUseBw   secAvailBw
```