Chapter 5

Configuring Physical Interface Properties

The software driver for each network media type sets reasonable default values for general interface properties, such as the interface's maximum transmission unit (MTU) size, receive and transmit leaky bucket properties, link operational mode, and clock source. To modify any of the default general interface properties, include the appropriate statements at the [edit interfaces *interface-name*] hierarchy level:

```
interfaces {
    interface-name {
        accounting-profile name;
        aggregated-ether-options {
            aggregated-ether-interface-options;
        }
        aggregated-sonet-options {
            aggregated-sonet-interface-options;
        }
        atm-options {
            atm-interface-options;
        }
        clocking clock-source;
        dce;
        description text;
        disable;
        ds0-options {
            ds0-interface-options;
        }
        e1-options {
            e1-interface-options;
        }
        e3-options {
            e3-interface-options;
        }
        encapsulation type;
        es-options {
            es-interface-options;
        }
        fastether-options {
            fastether-interface-options;
        }
        gigether-options {
            gigether-interface-options;
        }
```

```
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
   lmi-type (ansi | itu);
   n391dte number;
   n392dce number;
   n392dte number;
   n393dce number;
   n393dte number;
   t391dte seconds;
   t392dce seconds;
}
mac mac-address;
mlfr-uni-nni-bundle-options {
   mlfr-interface-options;
}
mtu bytes;
multiservice-options {
   boot-command filename;
   (core-dump | no-core-dump);
   (syslog | no-syslog);
}
no-gratuitous-arp-request;
no-keepalives;
no-partition {
   interface-type type;
}
partition partition-number oc-slice oc-slice-range interface-type type {
   timeslots time-slot-range;
}
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
   chap {
      access-profile name;
      local-name name;
      passive;
   }
}
receive-bucket {
   overflow (discard | tag);
   rate percentage;
   threshold bytes;
}
serial-options {
   serial-interface-options;
}
services-options {
   service-interface-options;
}
sonet-options {
   sonet-interface-options;
}
speed (10m | 100m);
stacked-vlan-tagging;
```

```
            t1-options {
                t1-interface-options;
            }
            t3-options {
                t3-interface-options;
            }
            traceoptions {
                flag flag <flag-modifier> <disable>;
            }
            transmit-bucket {
                overflow discard;
                rate percentage;
                threshold bytes;
            }
            (traps | no-traps);
            unit {
                logical-interface-statements;
            }
            vlan-tagging;
        }
    }
```

This chapter discusses configuration of the following physical interface properties:

Enabling or Disabling SNMP Notifications on Physical Interfaces on page 90

Disabling a Physical Interface on page 90

For information about interface-specific physical properties, see “Interface Types” on page 155.

Table 4 lists statements that you can use to configure physical interfaces.

Table 4: Statements for Physical Interface Properties

Statement	Interface Types	Usage Guidelines
802.3ad ae*x*	Aggregated Ethernet interfaces	“Configuring Ethernet Link Aggregation” on page 347 or “Configuring Aggregated Ethernet Interfaces” on page 392
access-profile *name*	Interfaces with Point-to-Point Protocol (PPP) encapsulation	“Configuring the PPP Challenge Handshake Authentication Protocol” on page 79
accounting-profile *name*	All	“Configuring Accounting for the Physical Interface” on page 85
acknowledge-retries *number*	Link services and voice services interfaces	“Configuring Link Services Acknowledgment Timers” on page 424
acknowledge-timer *milliseconds*	Link services and voice services interfaces	“Configuring Link Services Acknowledgment Timers” on page 424
action-red-differential-delay (disable-tx | remove-link)	Link services and voice services interfaces	“Configuring Link Services Differential Delay” on page 425
advertise-interval *milliseconds*	SONET/SDH interfaces	“Configuring APS and MSP” on page 522
aggregate	Gigabit Ethernet intelligent queuing (IQ) interfaces and Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	“Configuring Gigabit Ethernet Policers” on page 375
aggregate as*x*	Aggregated SONET/SDH interfaces	“Configuring Aggregated SONET/SDH Interfaces” on page 538
aggregated-ether-options	Aggregated Ethernet interfaces	“Configuring Aggregated Ethernet Interfaces” on page 392
aggregated-sonet-options	Aggregated SONET/SDH interfaces	“Configuring Aggregated SONET/SDH Interfaces” on page 538
aps	SONET/SDH interfaces	“Configuring APS and MSP” on page 522
atm-encapsulation (direct | plcp)	E3 and T3 traffic over Asynchronous Transfer Mode (ATM) interfaces	“Configuring E3 and T3 Parameters on ATM Interfaces” on page 234
atm-options	ATM1 and ATM2 IQ interfaces	“Configuring ATM Interfaces” on page 182
authentication-key *key*	SONET/SDH interfaces	“Configuring APS and MSP” on page 522
bandwidth-limit *bps*	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	“Configuring Gigabit Ethernet Policers” on page 375
bert-algorithm *algorithm*	E3, T1, T3, multichannel DS3, channelized interfaces (DS3, OC12, and STM1), and channelized IQ interfaces (E1 and DS3)	“Configuring BERT Properties” on page 86

Statement	Interface Types	Usage Guidelines
bert-error-rate *rate*	E1, E3, T1, T3, and channelized interfaces (DS3, OC3, OC12, and STM1)	"Configuring BERT Properties" on page 86
bert-period *seconds*	E1, E3, T1, T3, and channelized interfaces (DS3, OC12, and STM1)	"Configuring BERT Properties" on page 86
boot-command *filename*	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 89
boot-command *length*	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Configuring Gigabit Ethernet Policers" on page 375
buildout *value*	T1 interfaces	"Configuring the T1 Buildout" on page 547
buildout *feet*	E3 and T3 traffic over ATM interfaces	"Configuring E3 and T3 Parameters on ATM Interfaces" on page 234
burst-size-limit *bytes*	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Configuring Gigabit Ethernet Policers" on page 375
byte-encoding (nx64 | nx56)	DS0 and T1 interfaces	"Configuring T1 Byte Encoding" on page 548
bytes [*values*]	SONET/SDH interfaces	"Configuring SONET/SDH Header Byte Values" on page 512
cell-bundle-size *cells*	ATM2 IQ interfaces using ATM Layer 2 circuit cell-relay transport mode	"Configuring the Layer 2 Circuit Cell-Relay Cell Maximum" on page 208
(cbit-parity | no-cbit-parity)	T3 interfaces	"Disabling T3 C-Bit Parity Mode" on page 557
cbr *rate*	ATM interfaces	"Defining the ATM Traffic-Shaping Profile" on page 214
chap	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 79
classifier	Gigabit Ethernet IQ interfaces	"Specifying an Output Priority Map" on page 377
clock-rate *rate*	Serial interfaces (EIA-530 and V.35)	"Configuring the DTE Clock Rate" on page 500
clocking *clock-source*	ATM, DS0, E1, E3, SONET/SDH, T1, and T3 interfaces	"Configuring the Clock Source" on page 83
clocking-mode (dce | dte | loop)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Clocking Mode" on page 498
compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate *value*>	E3 and T3 interfaces	"Configuring the E3 CSU Compatibility Mode" on page 331 and "Configuring the T3 CSU Compatibility Mode" on page 558
control-leads	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Signal Handling" on page 501
control-polarity (positive | negative)	Serial interfaces (X.21)	"Configuring Serial Signal Polarities" on page 504
control-signal (assert | de-assert | normal)	Serial interfaces (X.21)	"Configuring the Serial Signal Handling" on page 501
(core-dump | no-core-dump)	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 89

Statement	Interface Types	Usage Guidelines
cts (ignore | normal | require)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 501
cts-polarity (positive | negative)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 504
dcd (ignore | normal | require)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 501
dcd-polarity (positive | negative)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 504
dce	Interfaces with Frame Relay encapsulation	"Configuring the Router as a DCE" on page 83
description *text*	All	"Adding an Interface Description to the Configuration" on page 66
disable	All	"Disabling a Physical Interface" on page 90
drop-timeout *milliseconds*	Multilink, link services, and voice services interfaces	"Configuring a Drop Timeout Period" on page 414
ds0-options	DS0 interfaces	"Channelized Interfaces Overview" on page 247
dsr (ignore | normal | require)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 501
dsr-polarity (positive | negative)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 504
dtr *signal-handling-option*	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 501
dtr-circuit (balanced | unbalanced)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial DTR Circuit" on page 504
dtr-polarity (positive | negative)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 504
e1-options	E1 interfaces	"Configuring E1 Interfaces" on page 321
e3-options	E3 interfaces	"Configuring E3 Interfaces" on page 329
encapsulation *type*	All interfaces, except loopback and multicast tunnel	"Configuring Interface Encapsulation" on page 73
encoding (nrz | nrzi)	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring Serial Line Encoding" on page 507
epd-threshold *cells*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
es-options	ES interfaces	"Configuring ES PIC Redundancy" on page 339
ethernet-policer-profile	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Configuring Gigabit Ethernet Policers" on page 375
ethernet-switch-profile	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Configuring Gigabit Ethernet Policers" on page 375, "Configuring MAC Address Filtering" on page 379, and "Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags" on page 382

Statement	Interface Types	Usage Guidelines
facility-override *facility-name*	Adaptive services interfaces	"Configuring Default System Log Properties" on page 162
fastether-options	Fast Ethernet interfaces	"Configuring Ethernet Physical Interface Properties" on page 344
fcs (32 | 16)	E1/E3, SONET/SDH, and T1/T3 interfaces	"Configuring the E1 Frame Checksum" on page 323, "Configuring the E3 Frame Checksum" on page 332, "Configuring the SONET/SDH Frame Checksum" on page 514, "Configuring the T1 Frame Checksum" on page 549, and "Configuring the T3 Frame Checksum" on page 560
(feac-loop-respond | no-feac-loop-respond)	T3 interfaces	"Configuring the T3 FEAC Response" on page 560
(flow-control | no-flow-control)	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	"Configuring Flow Control" on page 353
force	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
forwarding-class *class-name*	Gigabit Ethernet IQ and ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236 and "Specifying an Output Priority Map" on page 377
fragment-threshold *bytes*	Multilink, link services, and voice services interfaces	"Configuring a Fragmentation Threshold" on page 416
framing (g704 | g704-no-crc4 | g.751 | g.832 | unframed | sf | esf)	E1, E3, and T1 interfaces	"Configuring E3 and T3 Parameters on ATM Interfaces" on page 234, "Configuring E1 Framing" on page 324, and "Configuring T1 Framing" on page 550
gigether-options	Gigabit Ethernet interfaces	"Configuring Ethernet Physical Interface Properties" on page 344
(gratuitous-arp-reply | no-gratuitous-arp-reply)	Ethernet interfaces	"Configuring Gratuitous ARP" on page 354
hello-timer *milliseconds*	Link services and voice services interfaces	"Configuring Link Services Acknowledgment Timers" on page 424
high-plp-threshold *percent*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
hold-time *milliseconds*	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
hold-time up *milliseconds* down *milliseconds*	All interfaces, except aggregated SONET/SDH, generalized routing encapsulation (GRE) tunnel, and IP tunnel	"Damping Interface Transitions" on page 89 and "Configuring SONET/SDH Defect Triggers to Be Ignored" on page 518
host *hostname*	Adaptive services interfaces	"Configuring Default System Log Properties" on page 162
ieee802.1p premium [*values*]	Gigabit Ethernet IQ interfaces	"Specifying an Input Priority Map" on page 377
idle-cycle-flag *value*	E1, E3, T1, and T3 interfaces	"Configuring the E1 Idle Cycle Flag" on page 324, "Configuring the E3 Idle Cycle Flag" on page 333, "Configuring the T1 Idle Cycle Flag" on page 552, and "Configuring the T3 Idle Cycle Flag" on page 561

Statement	Interface Types	Usage Guidelines
ignore-all	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Signal Handling" on page 501
ilmi	ATM interfaces	"Configuring Communication with Directly Attached ATM Switches and Routers" on page 184
inactivity-timeout *seconds*	Adaptive services interfaces	"Configuring Default Timeout Settings" on page 162
indication (ignore | normal | require)	Serial interfaces (X.21)	"Configuring the Serial Signal Handling" on page 501
indication-polarity (positive | negative)	Serial interfaces (X.21)	"Configuring Serial Signal Polarities" on page 504
ingress-rate-limit *rate*	8-port, 12-port, and 48-port Fast Ethernet interfaces	"Configuring the Ingress Rate Limit" on page 355
input-priority-map	Gigabit Ethernet IQ interfaces	"Specifying an Input Priority Map" on page 377
interface-type *type*	Channelized IQ interfaces	"Channelized Interfaces Overview" on page 247
invert-data	DS0, E1, and T1 interfaces	"Configuring E1 Data Inversion" on page 324 and "Configuring T1 Data Inversion" on page 548
keepalives <down-count *number*> <interval *seconds*> <up-count *number*>	Aggregated SONET/SDH, DS0, E1, E3, SONET/SDH, T1, and T3 interfaces	"Configuring Keepalives" on page 82
lacp *mode*	Aggregated Ethernet interfaces	"Configuring Aggregated Ethernet LACP" on page 348
log-prefix *prefix-number*	Adaptive services interfaces	"Configuring Default System Log Properties" on page 162
line-encoding (ami | b8zs)	T1 interfaces	"Configuring T1 Line Encoding" on page 550
line-protocol *protocol*	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Line Protocol" on page 495
linear-red-profile *profile-name*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
linear-red-profiles *profile-name*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
link-mode *mode*	Management Ethernet (fxp0) and Fast Ethernet interfaces	"Configuring the Link Characteristics" on page 67
link-speed *speed*	Aggregated Ethernet and aggregated SONET/SDH interfaces	"Configuring Aggregated Ethernet Link Speed" on page 350 and "Configuring Aggregated SONET/SDH Link Speed" on page 540
lmi *lmi-options*	Interfaces with Frame Relay encapsulation	"Configuring Tunable Keepalives for Frame Relay LMI" on page 405 and "Configuring Link Services Keepalive Settings on Frame Relay LMI" on page 426
lmi-type (ansi | itu)	Link services interfaces and interfaces with Frame Relay encapsulation	"Configuring Frame Relay Keepalives" on page 404
local-name *name*	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 79

Statement	Interface Types	Usage Guidelines
lockout	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
(long-buildout | no-long-buildout)	T3 interfaces	"Configuring the T3 Line Buildout" on page 561
(loop-timing | no-loop-timing)	Channelized T3 interfaces	"Configuring the Channelized T3 Loop Timing" on page 562
loopback *mode*	DS0, E1, E3, T1, T3, Ethernet, SONET/SDH, and serial interfaces (EIA-530, V.35, and X.21)	"Configuring E1 Loopback Capability" on page 325, "Configuring E3 Loopback Capability" on page 333, "Configuring Ethernet Loopback Capability" on page 352, "Configuring Serial Loopback Capability" on page 505, "Configuring SONET/SDH Loopback Capability" on page 514, "Configuring T1 Loopback Capability" on page 551, and "Configuring T3 Loopback Capability" on page 562
(loopback | no-loopback)	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	"Configuring Ethernet Loopback Capability" on page 352
loss-priority (high | low)	Gigabit Ethernet IQ interfaces	"Specifying an Output Priority Map" on page 377
low-plp-max-threshold *percent*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
mac *mac-address*	Management Ethernet interface (fxp0)	"Configuring the MAC Address on the Management Ethernet Interface" on page 391
(mac-learn-enable | no-mac-learn-enable)	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Configuring MAC Address Filtering" on page 379
maximum-vcs *maximum-vcs*	ATM interfaces	"Configuring the Maximum Number of ATM1 VCs on a VP" on page 193
minimum-links *number*	Multilink, link services, and voice services interfaces	"Configuring Link Services and Multilink Interfaces" on page 409
mlfr-uni-nni-bundle-options *bundle-options*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
mpls	ATM and SONET/SDH interfaces in passive monitoring mode	"Removing MPLS Labels from Incoming Packets" on page 186 and page 531
mrru *bytes*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
mtu *bytes*	All interfaces, except management Ethernet (fxp0), loopback, multilink, and multicast tunnel	"Configuring the Media MTU" on page 67
multiservice-options	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 89
n391 *number*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
n392 *number*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
n393 *number*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
neighbor *address*	SONET/SDH interfaces	"Configuring APS and MSP" on page 522

Statement	Interface Types	Usage Guidelines
no-gratuitous-arp-request	Ethernet interfaces	"Configuring Gratuitous ARP" on page 354
no-keepalives	Interfaces with PPP, Frame Relay, or Cisco High-level Data Link Control (HDLC) encapsulation	"Configuring Keepalives" on page 82
no-partition	Channelized IQ interfaces	"Channelized Interfaces Overview" on page 247
oam-liveness	ATM interfaces	"Configuring the OAM F4 Cell Flows" on page 210
oam-period (disable | *seconds*)	ATM interfaces	"Configuring the OAM F4 Cell Flows" on page 210
oc-slice *oc-slice-range*	Channelized OC12 IQ interfaces	"Configuring Channelized OC12 Interfaces" on page 273
open-timeout *seconds*	Adaptive services interfaces	"Configuring Default Timeout Settings" on page 162
output-priority-map	Gigabit Ethernet IQ interfaces	"Specifying an Output Priority Map" on page 377
overflow (discard | tag)	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	"Configuring Receive and Transmit Leaky Bucket Properties" on page 84 and page 533
paired-group *group-name*	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
partition *partition-number*	Channelized IQ interfaces	"Channelized Interfaces Overview" on page 247
passive	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 79
passive-monitor-mode	SONET/SDH interfaces	"Enabling Passive Monitoring on SONET/SDH Interfaces" on page 530
path-trace *trace-string*	SONET/SDH interfaces	"Configuring the SONET/SDH Path Trace Identifier" on page 516
(payload-scrambler | no-payload-scrambler)	E3, SONET/SDH, and T3 interfaces	"Configuring E3 and T3 Parameters on ATM Interfaces" on page 234, "Configuring E3 HDLC Payload Scrambling" on page 335, "Configuring SONET/SDH HDLC Payload Scrambling" on page 516, "Configuring T3 HDLC Payload Scrambling" on page 564, and "Examples: Configuring T3 Interfaces" on page 565
per-unit-scheduler	IQ interfaces	"Associating a Scheduler Map with a DLCI or VLAN" on page 846
pic-type	ATM2 IQ interfaces	"Configuring the ATM PIC Type" on page 188
plp1 *cells*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
plp-to-clp	ATM2 IQ interfaces	"Enabling the PLP Setting to Be Copied to the CLP Bit" on page 240
policer *cos-policer-name*	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Configuring Gigabit Ethernet Policers" on page 375

Statement	Interface Types	Usage Guidelines
pop-all-labels	ATM and SONET/SDH interfaces in passive monitoring mode	"Removing MPLS Labels from Incoming Packets" on page 186 and page 531
ppp-options	Interfaces with PPP encapsulation	"Configuring the PPP Challenge Handshake Authentication Protocol" on page 79
premium	Gigabit Ethernet IQ interfaces	"Configuring Gigabit Ethernet Policers" on page 375 and "Specifying an Output Priority Map" on page 377
promiscuous-mode vpi *vpi-identifier*	ATM2 IQ interfaces	"Configuring ATM Cell-Relay Promiscuous Mode" on page 189
protect-circuit *group-name*	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
queue-depth *cells*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
queue-length *number*	ATM1 interfaces	"Configuring the ATM1 Queue Length" on page 221
rate *percentage*	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	"Configuring Receive and Transmit Leaky Bucket Properties" on page 84 and page 533
receive-bucket	All interfaces, except ATM, Fast Ethernet, and Gigabit Ethernet	"Configuring Receive and Transmit Leaky Bucket Properties" on page 84
red-differential-delay *milliseconds*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
remote-loopback-respond	T1 interfaces	"Configuring the T1 Remote Loopback Response" on page 549
request	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
required-depth *number*	ATM and SONET/SDH interfaces in passive monitoring mode	"Removing MPLS Labels from Incoming Packets" on page 186 and page 531
revert-time *seconds*	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
rfc-2615	SONET/SDH interfaces	"Configuring SONET/SDH RFC 2615 Support" on page 517
rts (assert | de-assert | normal)	Serial interfaces (EIA-530 and V.35)	"Configuring the Serial Signal Handling" on page 501
rts-polarity (positive | negative)	Serial interfaces (EIA-530 and V.35)	"Configuring Serial Signal Polarities" on page 504
rtvbr peak *rate* sustained *rate* burst *length*	ATM interfaces	"Configuring ATM2 IQ Real-Time VBR" on page 216
scheduler-maps *map-name*	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
serial-options	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring Serial Interfaces" on page 493
services-options	Services interfaces	*JUNOS Services Interf aces Confi guration Guide*
[services *priority-level*]	Adaptive services interfaces	"Configuring Default System Log Properties" on page 162
shaping	ATM interfaces	"Defining the ATM Traffic-Shaping Profile" on page 214
sonet-options	SONET/SDH interfaces	"Configuring SONET/SDH Physical Interface Properties" on page 510

Statement	Interface Types	Usage Guidelines
source-address-filter *mac-address*	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	"Enabling Ethernet MAC Address Filtering" on page 351
(source-filtering | no-source-filtering)	Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ, and Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Enabling Ethernet MAC Address Filtering" on page 351
speed (10m | 100m)	Management Ethernet interface (fxp0) and 12-port and 48-port Fast Ethernet interfaces	"Configuring the Interface Speed" on page 81
stacked-vlan-tagging	Gigabit Ethernet IQ interfaces	"Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags" on page 382
start-end-flag (shared | filler)	DS0, E1, E3, T1, and T3 interfaces	"Configuring E1 Start End Flags" on page 326, "Configuring the E3 Start End Flags" on page 335, "Configuring T1 Start End Flags" on page 553, and "Configuring the T3 Start End Flags" on page 565
switching-mode (bidirectional | unidirectional)	Unchannelized OC3, OC12, and OC48 SONET/SDH interfaces on T-series platforms	"Configuring Switching Between the Working and Protect Circuits" on page 525
syslog	Adaptive services interfaces	"Configuring Default System Log Properties" on page 162
(syslog | no-syslog)	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 89
t1-options	T1 interfaces	"Configuring T1 Interfaces" on page 545
t3-options	T3 interfaces	"Configuring T3 Interfaces" on page 555
t391 *seconds*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
t392 *number*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
tag-protocol-id [*tpids*]	Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform)	"Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags" on page 382
threshold *bytes*	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	"Configuring Receive and Transmit Leaky Bucket Properties" on page 84 and page 533
timeslots *time-slot-range*	Channelized T1 IQ and channelized E1 IQ interfaces	"Channelized Interfaces Overview" on page 247
tm (ignore | normal | require)	Serial interfaces (EIA-530)	"Configuring the Serial Signal Handling" on page 501
tm-polarity (positive | negative)	Serial interfaces (EIA-530)	"Configuring Serial Signal Polarities" on page 504
traceoptions	All	"Tracing Operations of an Individual Router Interface" on page 88
transmit-bucket	All interfaces, except ATM, Fast Ethernet, and Gigabit Ethernet	"Configuring Receive and Transmit Leaky Bucket Properties" on page 84
transmit-clock invert	Serial interfaces (EIA-530, V.35, and X.21)	"Configuring the Serial Clocking Mode" on page 498

Statement	Interface Types	Usage Guidelines
(traps | no-traps)	All	"Enabling or Disabling SNMP Notifications on Physical Interfaces" on page 90
trigger *defect* ignore	ATM over SONET/SDH and SONET/SDH interfaces	"Configuring SONET/SDH Defect Triggers to Be Ignored" on page 518
(unframed | no-unframed)	E3 IQ interfaces	"Configuring E3 IQ Unframed Mode" on page 336
vbr peak *rate* sustained *rate* burst *length*	ATM interfaces	"Defining the ATM Traffic-Shaping Profile" on page 214
vpi *vpi-identifier*	ATM interfaces	"Configuring ATM Cell-Relay Promiscuous Mode" on page 189 and "Configuring the Maximum Number of ATM1 VCs on a VP" on page 193
vc-cos-mode (alternate | strict)	ATM2 interfaces	"Configuring ATM2 IQ VC Tunnel CoS Components" on page 236
vlan-tagging	Fast Ethernet and Gigabit Ethernet interfaces	"Configuring 802.1Q VLANs" on page 355
vtmapping	Channelized STM1 interfaces	"Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces" on page 297
working-circuit *group-name*	SONET/SDH interfaces	"Configuring APS and MSP" on page 522
yellow-differential-delay *milliseconds*	Link services and voice services interfaces	"Configuring Link Services Physical Interface Properties" on page 422
(z0-increment | no-z0-increment)	SONET/SDH interfaces	"Configuring an Incrementing STM ID" on page 513

You specify aggregated interfaces by assigning a number for the aggregated interface. For aggregated Ethernet interfaces, configure ae*x* as in the following example:

```
[edit interfaces]
ae0 {
...
}
```

For aggregated SONET/SDH interfaces, configure as*x* as in the following example:

```
[edit interfaces]
as0 {
...
}
```

The maximum number of aggregated interfaces is 16, and the assigned number can be from 0 through 15. You should not mix SONET and SDH mode on the same aggregated interface.

NOTE: SONET/SDH aggregation is proprietary to the JUNOS software and might not work with other software.

For aggregated Ethernet interfaces, you must include the vlan-tagging statement at the [edit interfaces ae*x*] hierarchy level to complete the association.

For more information, see “Configuring Aggregated Ethernet Interfaces” on page 392 and “Configuring Aggregated SONET/SDH Interfaces” on page 538.

Adding an Interface Description to the Configuration

You can include a text description of each physical interface in the configuration file. Any descriptive text you include is displayed in the output of the show interfaces commands, and is also exposed in the ifAlias Management Information Base (MIB) object. It has no impact on the interface’s configuration. To add a text description, include the description statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
description text;
```

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.

For information about describing logical units, see “Adding a Logical Unit Description to the Configuration” on page 100.

Example: Adding an Interface Description to the Configuration

Add a description to a SONET/SDH interface:

```
[edit interfaces so-1/1/0]
user@host# set description "BB: phl01 P12/0/0 - local wire"
[edit interfaces so-1/1/0]
user@host# commit
[edit interfaces so-1/1/0]
user@host# exit configuration-mode
cli> show interfaces so-1/1/0
so-1/1/0 {
    physical-interface index 9 snmp-ifindex 10;
    enabled physical-link up;
    description "BB: phl01 P12/0/0 - local wire";
    encapsulation cisco-hdlc;
...
```

Configuring the Link Characteristics

By default, the routing platform's management Ethernet interface, fxp0, autonegotiates whether to operate in full-duplex or half-duplex mode. Fast Ethernet interfaces can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet, the link partner must also be set to full duplex.

To explicitly configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the link-mode statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
link-mode (full-duplex | half-duplex);
```

Configuring the Media MTU

The default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO). Table 5 through Table 10 on page 70 list the media and protocol MTU sizes by interface type, and Table 11 on page 71 lists the encapsulation overhead by encapsulation type.

Table 5: Media MTU Sizes by Interface Type for M5, M7i, M10, M10i, M20, and M40 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	9192	1500
E3/T3	4474	9192	4470
Fast Ethernet	1514	9192 (4-port) 1532 (8-port) 1532 (12-port)	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
Serial	1504	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 6: Media MTU Sizes by Interface Type for M40e Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
Fast Ethernet	1514	4500	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 4500 (4-port)	1500 (IPv4) 1497 (ISO)
Serial	1504	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 4500 (4-port OC3) 4500 (4-port OC3c) 4500 (1-port OC12) 4500 (4-port OC12) 4500 (4-port OC12c) 4500 (1-port OC48) 9192 (2-port OC3) 9192 (2-port OC3c) 9192 (1-port OC12c) 9192 (1-port OC48c) 4500 (1-port OC192) 9192 (1-port OC192c)	4470

Table 7: Media MTU Sizes by Interface Type for M160 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
Fast Ethernet	1514	4500	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 4500 (4-port)	1500 (IPv4) 1497 (ISO)
Serial	1504	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (1- or 2-port) 4500 (4-port)	4470

Table 8: Media MTU Sizes by Interface Type for M320 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
Channelized DS3 IQ	4471	4500	4470
Channelized E1 IQ	1504	4500	1500
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
E1	1504	4500	1500
E3 IQ	4471	4500	4470
Fast Ethernet	1514	4500 (4-port) 1532 (12- and 48-port)	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470
T1	1504	4500	1500

Table 9: Media MTU Sizes by Interface Type for T320 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM	4482	9192	4470
ATM2 IQ	4482	9192	4470
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
Fast Ethernet	1514	4500 (4-port) 1532 (12- and 48-port)	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 10: Media MTU Sizes by Interface Type for T640 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
48-port Fast Ethernet	1514	1532	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 11: Encapsulation Overhead by Encapsulation Type

Interface Encapsulation	Encapsulation Overhead (Bytes)
802.1Q/Ethernet 802.3	21
802.1Q/Ethernet Subnetwork Access Protocol (SNAP)	26
802.1Q/Ethernet version 2	18
ATM Cell Relay	4
ATM permanent virtual connection (PVC)	12
Cisco HDLC	4
Ethernet 802.3	17
Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS)	4
Ethernet over ATM	32
Ethernet SNAP	22
Ethernet translational cross-connect (TCC)	18
Ethernet version 2	14
Extended virtual local area network (VLAN) CCC and VPLS	4
Extended VLAN TCC	22
Frame Relay	4
PPP	4
VLAN CCC	4
VLAN VPLS	4
VLAN TCC	22

The default media MTU is calculated as follows:

Default media MTU = Default IP MTU + encapsulation overhead

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.

NOTE: The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet interface is specified as 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

A SONET/SDH interface operating in concatenated mode has a “c” added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

If you do not configure an MPLS MTU, the JUNOS software derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, so 12 bytes.

In other words, the formula used to determine the MPLS MTU is physical interface MTU - encapsulation overhead - 12.

If you configure an MTU value by including the mtu statement at the [edit interfaces *interface-name* unit *logical-unit-number* family mpls] hierarchy level, the configured value is used.

For information about configuring the encapsulation on an interface, see “Configuring Interface Encapsulation” on page 73.

To modify the default media MTU size for a physical interface, include the mtu statement at the [edit interfaces *interface-name*] hierarchy level:

[edit interfaces *interface-name*]
mtu *bytes*;

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. You configure the protocol MTU by including the mtu statement at the following hierarchy levels:

[edit interfaces *interface-name* unit *logical-unit-number* family *family*]

[edit logical-routers *logical-router-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

For more information, see “Setting the Protocol MTU” on page 117.

Configuring Interface Encapsulation

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types. For more information about logical interface encapsulation, see "Configuring the Encapsulation on a Logical Interface" on page 105.

This section is organized as follows:

Configuring the Encapsulation on a Physical Interface

By default, PPP is the encapsulation type for physical interfaces. To configure the encapsulation on a physical interface, include the encapsulation statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc |
cisco-hdlc-tcc | ethernet-ccc | ethernet-over-atm | ethernet-tcc | ethernet-vpls |
extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-ccc |
extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services |
flexible-frame-relay | frame-relay | frame-relay-ccc | frame-relay-port-ccc |
frame-relay-tcc | multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc |
vlan-vpls);
```

The physical interface encapsulation can be one of the following:

- ATM CCC Cell Relay—Connects two remote virtual circuits or ATM physical interfaces with a label-switched path (LSP). Traffic on the circuit is ATM cells.

  You can configure an ATM1 Physical Interface Card (PIC) to use cell-relay accumulation mode (CAM). In this mode, the incoming cells (1 to 8 cells) are packaged into a single packet and forwarded to the LSP. Cell-relay accumulation mode is not supported on ATM2 PICs. You configure CAM as shown in the following example:

  ```
  [edit chassis]
  fpc 1 {
      pic 0 {
          atm-cell-relay-accumulation;
      }
  }
  ```

  For more information, see the *JUNOS System Basics Configuration Guide*.

- ATM PVC—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routing platforms using the Label Distribution Protocol (LDP).

- Cisco HDLC—E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:

  - CCC version (cisco-hdlc-ccc)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the ccc family only.

  - TCC version (cisco-hdlc-tcc)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

- Ethernet over ATM—As defined in RFC 1483, this encapsulation type allows ATM interfaces to connect to devices that support only bridged-mode protocol data units (BPDUs). The JUNOS software does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet logical link control (LLC)/SNAP frames with IP or Address Resolution Protocol (ARP) in the payload, and drops the rest. For packets destined to the Ethernet local area network (LAN), a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and media access control (MAC) header, and the packet is forwarded to the ATM interface.

Ethernet cross-connect—Ethernet interfaces without VLAN tagging can use Ethernet CCC encapsulation. Two related versions are supported:

CCC version (ethernet-ccc)—Ethernet interfaces with standard Tag Protocol ID (TPID) tagging can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the ccc family only.

TCC version (ethernet-tcc)—Similar to CCC, but used for circuits with different media on either side of the connection. One-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs can use Ethernet TCC encapsulation.

VLAN CCC (vlan-ccc)—Ethernet interfaces with VLAN tagging enabled can use VLAN CCC encapsulation. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the ccc family only.

Extended VLAN cross-connect—Gigabit Ethernet interfaces with VLAN 802.1Q tagging enabled can use extended VLAN cross-connect encapsulation. (Ethernet interfaces with standard TPID tagging can use VLAN CCC encapsulation.) Two related versions of extended VLAN cross-connect are supported:

CCC version (extended-vlan-ccc)—Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. Extended VLAN CCC is not supported on 4-port Gigabit Ethernet PICs. When you use this encapsulation type, you can configure the ccc family only.

TCC version (extended-vlan-tcc)—Similar to CCC, but used for circuits with different media on either side of the connection. One-port Gigabit Ethernet, 2-port Gigabit Ethernet, and 4-port Fast Ethernet PICs can use Extended Ethernet TCC encapsulation.

Ethernet VPLS (ethernet-vpls)—Ethernet interfaces with VPLS enabled can use Ethernet VPLS encapsulation. For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.

Ethernet VLAN VPLS (vlan-vpls)—Ethernet interfaces with VLAN tagging and VPLS enabled can use Ethernet VLAN VPLS encapsulation. For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.

Extended VLAN VPLS (extended-vlan-vpls)—Ethernet interfaces with VLAN 802.1Q tagging and VPLS enabled can use Ethernet Extended VLAN VPLS encapsulation. (Ethernet interfaces with standard TPID tagging can use Ethernet VLAN VPLS encapsulation.) Extended Ethernet VLAN VPLS encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.

Flexible Ethernet services (flexible-ethernet-services)—Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform) can use flexible Ethernet services encapsulation. Aggregated Ethernet bundles cannot use this encapsulation type. You use this encapsulation type when you want to configure multiple per-unit Ethernet encapsulations. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

Flexible Frame Relay (flexible-frame-relay)—IQ interfaces can use flexible Frame Relay encapsulation. You use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any data-link connection identifier (DLCI) value from 1 through 1022.

Frame Relay (frame-relay)—Defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation. Five related versions are supported:

CCC version (frame-relay-ccc)—The same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. This numbering restriction does not apply to IQ interfaces. The logical interface must also have frame-relay-ccc encapsulation. When you use this encapsulation type, you can configure the ccc family only.

TCC version (frame-relay-tcc)—Similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

Extended CCC version (extended-frame-relay-ccc)—This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. The logical interface must have frame-relay-ccc encapsulation. When you use this encapsulation type, you can configure the ccc family only.

Extended TCC version (extended-frame-relay-tcc)—Similar to extended Frame Relay CCC, this encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC, which is used for circuits with different media on either side of the connection.

Port CCC version (frame-relay-port-ccc)—Defined in the IETF document *Frame Relay Encapsulation over Pseudo-Wires* (expired December 2002). This encapsulation type allows you to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the ccc family only.

Multilink Frame Relay (MLFR) UNI and NNI (multilink-frame-relay-uni-nni)—Link services and voice services interfaces functioning as FRF.16 bundles can use multilink Frame Relay UNI NNI encapsulation. This encapsulation is also used on link services and voice services interfaces' constituent T1, E1, or *N*xDS0 interfaces.

PPP—Defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation. Two related versions are supported:

Circuit cross-connect (CCC) version (ppp-ccc)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the ccc family only.

Translational cross-connect (TCC) version (ppp-tcc)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.

Encapsulation Capabilities

When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one unit statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.

Ethernet CCC encapsulation for Ethernet interfaces with standard TPID tagging requires that the physical interface have only a single logical interface. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For Ethernet interfaces in VLAN mode, VLAN IDs are applicable as follows:

VLAN ID 0 is reserved for tagging the priority of frames.

For encapsulation type vlan-ccc, VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN CCCs.

For encapsulation type vlan-vpls, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLANs.

For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i platform), you can configure flexible Ethernet services encapsulation on the physical interface. For interfaces with flexible-ethernet-services encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511 are not reserved.

For encapsulation types extended-vlan-ccc and extended-vlan-vpls, all VLAN IDs are valid.

The upper limits for configurable VLAN IDs vary by interface type. For more information, see "Configuring 802.1Q VLANs" on page 355.

When you configure a TCC encapsulation, some modifications are needed to handle VPN connections over unlike Layer 2 and Layer 2.5 links and terminate the Layer 2 and Layer 2.5 protocol locally. The routing platform performs the following media-specific changes:

PPP TCC—Both Link Control Protocol (LCP) and Network Control Protocol (NCP) are terminated on the routing platform. Internet Protocol Control Protocol (IPCP) IP address negotiation is not supported. The JUNOS software strips all PPP encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to PPP encapsulation.

Cisco HDLC TCC—Keepalive processing is terminated on the routing platform. The JUNOS software strips all Cisco HDLC encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to Cisco HDLC encapsulation.

Frame Relay TCC—All Local Management Interface (LMI) processing is terminated on the routing platform. The JUNOS software strips all Frame Relay encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to Frame Relay encapsulation.

ATM—Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) processing is terminated at the routing platform. Cell relay is not supported. The JUNOS software strips all ATM encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to ATM encapsulation.

Example: Configuring the Encapsulation on a Physical Interface

Configure PPP encapsulation on a SONET/SDH interface. The second and third family statements allow Intermediate System-to-Intermediate System (IS-IS) and MPLS to run on the interface.

```
[edit interfaces]
so-7/0/0 {
    encapsulation ppp;
    unit 0 {
        point-to-point;
        family inet {
            address 192.168.1.113/32 {
                destination 192.168.1.114;
            }
        }
        family iso;
        family mpls;
    }
}
```

Configuring the PPP Challenge Handshake Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*. When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges. To enable CHAP, you must create an access profile, and you must configure the interfaces to use CHAP.

To configure a CHAP access profile, include the profile statement and specify a profile name at the [edit access] hierarchy level:

```
[edit access]
profile profile-name {
    client name chap-secret data;
}
```

For more information about configuring access profiles, see the *JUNOS System Basics Configuration Guide*.

When you configure an interface to use CHAP, you must assign an access profile to the interface. When an interface receives CHAP challenges and responses, the access profile in the packet is used to look up the shared secret, as defined in RFC 1994.

To configure PPP CHAP on an interface with PPP encapsulation, include the chap statement at the [edit interfaces *interface-name* ppp-options] hierarchy level:

```
[edit interfaces interface-name ppp-options]
chap {
    access-profile name;
    local-name name;
    passive;
}
```

On each interface with PPP encapsulation, you can configure the following PPP CHAP properties:

- Assigning an Access Profile to an Interface on page 80
- Configuring the Local Name on page 80
- Configuring Passive Mode on page 80

When you configure PPP over ATM or Multilink PPP over ATM encapsulation, you can enable CHAP on the logical interface. For more information, see "Configuring PPP over ATM2 Encapsulation" on page 230.

Assigning an Access Profile to an Interface

To assign an access profile to an interface, include the access-profile statement at the [edit interfaces *interface-name* ppp-options chap] hierarchy level:

[edit interfaces *interface-name* ppp-options chap]
access-profile *name*;

You must include the access-profile statement when you configure the CHAP authentication method. If an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped.

Configuring the Local Name

By default, when CHAP is enabled on an interface, the interface uses the routing platform's system hostname as the name sent in CHAP challenge and response packets.

To configure the name the interface uses in CHAP challenge and response packets, include the local-name statement at the [edit interfaces *interface-name* ppp-options chap] hierarchy level:

[edit interfaces *interface-name* ppp-options chap]
local-name *name*;

Configuring Passive Mode

By default, when CHAP is enabled on an interface, the interface always challenges its peer and responds to challenges from its peer.

You can configure the interface not to challenge its peer, and only respond when challenged. To configure the interface not to challenge its peer, include the passive statement at the [edit interfaces *interface-name* ppp-options chap] hierarchy level:

[edit interfaces *interface-name* ppp-options chap]
passive;

Example: Configuring the PPP Challenge Handshake Authentication Protocol

Configure CHAP:

```
[edit access]
    profile pe-A-ppp-clients;
      client cpe-1 chap-secret "$1$dQYsZ$B5ojUeUjDsUo.yKwcCZ0";
           # SECRET-DATA
      client cpe-2 chap-secret "$1$kdAsfaDAfkdjDsASxfafdKdFKJ";
           # SECRET-DATA
    }
}

[edit interfaces so-1/2/0]
encapsulation ppp;
ppp-options {
    chap {
      access-profile pe-A-ppp-clients;
      local-name "pe-A-so-1/1/1";
    }
}

[edit interfaces so-1/1/2]
encapsulation ppp;
ppp-options {
    chap {
      access-profile pe-A-ppp-clients;
      local-name “pe-A-so-1/1/2";
    }
}
```

Configuring the Interface Speed

By default, the routing platform’s management Ethernet interface, fxp0, autonegotiates whether to operate at 10 megabits per second (Mbps) or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement in the [edit chassis] configuration hierarchy, as described in the *JUNOS Sy stem Basics Configur ation Guide*).

To configure the management Ethernet interface to operate at 10 Mbps or 100 Mbps, include the speed statement at the [edit interfaces fxp0] hierarchy level:

```
[edit interfaces fxp0]
speed (10m | 100m);
```

Configuring Keepalives

By default, physical interfaces configured with Cisco HDLC or PPP encapsulation send keepalive packets at 10-second intervals. The Frame Relay term for keepalives is LMI packets; the JUNOS software supports both ANSI T1.617 Annex D LMIs and ITU Q933 Annex A LMIs. On ATM networks, OAM cells perform the same function. You configure OAM cells at the logical interface level; for more information, see "Defining the ATM OAM F5 Loopback Cell Period" on page 225.

To disable the sending of keepalives on a physical interface, include the no-keepalives statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
no-keepalives;
```

When you configure PPP over ATM or Multilink PPP over ATM encapsulation, you can enable or disable keepalives on the logical interface. For more information, see "Configuring PPP over ATM2 Encapsulation" on page 230.

To explicitly enable the sending of keepalives on a physical interface, include the keepalives statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
keepalives;
```

To change one or more of the default keepalive values, include the appropriate option at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
keepalives <interval seconds> <down-count number> <up-count number>;
```

On interfaces configured with Cisco HDLC or PPP encapsulation, you can include the following three keepalive statements; note that Frame Relay encapsulation is not affected by these statements:

- interval *seconds*—The time in seconds between successive keepalive requests. The range is from 1 second through 32767 seconds, with a default of 10 seconds.

- down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is from 1 through 255, with a default of 3.

- up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is from 1 through 255, with a default of 1.

For information about Frame Relay keepalive settings, see "Configuring Frame Relay Keepalives" on page 404.

Configuring the Clock Source

For interfaces such as SONET/SDH that can use different clock sources, you can configure the source of the transmit clock on each interface. The source can be internal (also called line timing or normal timing) or external (also called loop timing). The default source is internal, which means that each interface uses the routing platform's internal stratum 3 clock.

For T3 channels on a channelized OC12 interface, T1 channels on a channelized T3 interface, and DS0 channels on a channelized E1 interface, the clocking statement is supported only for channel 0; it is ignored if included in the configuration of other channels. The clock source configured for channel 0 applies to all channels on the channelized OC12, channelized DS3, and channelized E1 interfaces. The individual DS3, DS1, and DS0 channels use a gapped 45-MHz clock as the transmit clock. For more information, see "Clock Sources on Channelized Interfaces" on page 250.

NOTE: On channelized STM1 interfaces, you should configure the clock source at one side of the connection to be internal (the default JUNOS configuration) and configure the other side of the connection to be external.

To configure loop timing on an interface, include the clocking external statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
clocking external;
```

To explicitly configure line timing on an interface, include the clocking internal statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
clocking internal;
```

Configuring the Router as a DCE

By default, when you configure an interface with Frame Relay encapsulation, the routing platform is assumed to be data terminal equipment (DTE). That is, the routing platform is assumed to be at a terminal point on the network. To configure the routing platform to be data circuit-terminating equipment (DCE), include the dce statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
dce;
```

When you configure the routing platform to be a DCE, keepalives are disabled by default.

For back-to-back Frame Relay connections, either disable the sending of keepalives on both sides of the connection, or configure one side of the connection as a DTE (the default JUNOS configuration) and the other as a DCE.

Configuring Receive and Transmit Leaky Bucket Properties

Congestion control is particularly difficult in high-speed networks with high volumes of traffic. When congestion occurs in such a network, it is usually too late to react. You can avoid congestion by regulating the flow of packets into your network. Smoother flows prevent bursts of packets from arriving at (or being transmitted from) the same interface and causing congestion.

For all interface types except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ, you can configure leaky bucket properties, which allow you to limit the amount of traffic received on and transmitted by a particular interface. You effectively specify what percentage of the interface's total capacity can be used to receive or transmit packets. You might want to set leaky bucket properties to limit the traffic flow from a link that is known to transmit high volumes of traffic.

NOTE: Instead of configuring leaky bucket properties, you can limit traffic flow by configuring policers. Policers work on all interfaces. For more information, see "Applying Policers" on page 121 and the *JUNOS Policy Framework Configuration Guide*.

The leaky bucket is used at the host-network interface to allow packets into the network at a constant rate. Packets might be generated in a bursty manner, but after they pass through the leaky bucket, they enter the network evenly spaced. In some cases, you might want to allow short bursts of packets to enter the network without smoothing them out. By controlling the number of packets that can accumulate in the bucket, the threshold property controls burstiness. The maximum number of packets entering the network in *t* time units is threshold + rate * *t*.

By default, leaky buckets are disabled, and the interface can receive and transmit packets at the maximum line rate.

To configure leaky bucket properties, include one or both of the receive-bucket and transmit-bucket statements at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
```

In the rate statement, specify the percentage of the interface line rate that is available to receive or transmit packets. The percentage can be a value from 0 (none of the interface line rate is available) to 100 (the maximum interface line rate is available). For example, when you set the line rate to 33, the interface receives or transmits at one third of the maximum line rate.

In the threshold statement, specify the bucket threshold, which controls the burstiness of the leaky bucket mechanism. The larger the value, the more bursty the traffic, which means that over a very short amount of time the interface can receive or transmit close to line rate, but the average over a longer time is at the configured bucket rate. The threshold can be a value from 0 through 16,777,215 bytes. For ease of entry, you can enter *number* either as a complete decimal number or as a decimal number followed by the abbreviation k (1,000) or m (1,000,000). For example, the entry threshold 2m corresponds to a threshold of 2,000,000 bytes.

In the overflow statement, specify how to handle packets that exceed the threshold:

- tag (receive bucket only)—Tag, count, and process received packets that exceed the threshold.
- discard—Discard received packets that exceed the threshold. No counting is done.

Configuring Accounting for the Physical Interface

Juniper Networks routing platforms can collect various kinds of data about traffic passing through the routing platform. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the routing platform retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the [edit accounting-options] hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the interface-profile statement at the [edit accounting-options] hierarchy level. You configure filter profiles by including the filter-profile statement at the [edit accounting-options] hierarchy level. For more information, see the *JUNOS Network Management Configuration Guide*.

You apply filter profiles by including the accounting-profile statement at the [edit firewall filter *filter-name*] and [edit firewall family *family* filter *filter-name*] hierarchy levels. For more information, see the *JUNOS Policy Framework Configuration Guide*.

Applying an Accounting Profile to the Physical Interface

To enable accounting on an interface, include the accounting-profile statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
accounting-profile name;
```

You can also reference profiles by logical unit; for more information, see “Configuring Accounting for the Logical Interface” on page 102.

Example: Applying an Accounting Profile to the Physical Interface

Configure an accounting profile for an interface and apply it to a physical interface:

```
[edit]
accounting-options {
    file if_stats {
        size 4m files 10 transfer-interval 15;
        archive-sites {
            "ftp://login:password@host/path";
        }
    }
    interface-profile if_profile {
        interval 15;
        file if_stats {
            fields {
                input-bytes;
                output-bytes;
                input-packets;
                output-packets;
                input-errors;
                output-errors;
            }
        }
    }
}

[edit interfaces ge-1/0/1]
accounting-profile if_profile;
```

Configuring BERT Properties

Testing interfaces allows you to troubleshoot problems and check the quality of links. You can configure any of the following interfaces to execute a bit error rate test (BERT) when the interface receives a request to run this test: E1, E3, T1, T3, the channelized DS3, OC3, OC12, and STM1 interfaces, and the channelized DS3 IQ, E1, and OC12 IQ interfaces. On all specified interface types, you set the duration of the test and the error rate to include in the bit stream by including the bert-period and bert-error-rate statements at the [edit interfaces *interface-name interface*-options] hierarchy level:

```
[edit interfaces interface-name interface-options]
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. The algorithm for the E1 BERT procedure is pseudo-2e15-o151 (pattern is $2^{15}–1$, as defined in the CCITT/ITU O.151 standard). On T1, E3, T3, *N*xDS0, and channelized E1 and T3 IQ interfaces, you can also select the pattern to send in the bit stream by including the bert-algorithm statement at the [edit interfaces *interface-name interface*-options] hierarchy level:

```
[edit interfaces interface-name interface-options]
bert-algorithm algorithm;
```

For a list of supported algorithms, enter a ? after the bert-algorithm statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152     Pattern is 2^11 -1 (per O.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per O.152 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e20-o153     Pattern is 2^20 - 1 (per O.153 standard)
...
```

For specific hierarchy information, see individual interface types. For information about running the BERT procedure, see the *JUNOS Network and Services Interfaces Command Reference*.

Table 12 on page 87 shows the BERT capabilities for various interface types.

Table 12: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	Single port at a time Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	Single port at a time
Channelized OC12	N/A	Yes (channel 0–11)	Single channel at a time Limited algorithms No bit count
Channelized STM1	Yes (channel 0–62)	N/A	Multiple channels Only one algorithm No error insert No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	Multiple ports and channels Limited algorithms for T1 No error insert for T1 No bit count for T1

For information about BERT capabilities on channelized IQ interfaces, see “Channelized IQ Interface Properties” on page 252.

To exchange BERT patterns between a local routing platform and a remote routing platform, you include the loopback remote statement in the interface configuration at the remote end of the link. From the local routing platform, you issue the test interface command.

For more information about configuring loopbacks, see “Configuring E1 Loopback Capability” on page 325, “Configuring E3 Loopback Capability” on page 333, “Configuring SONET/SDH Loopback Capability” on page 514, “Configuring T1 Loopback Capability” on page 551, and “Configuring T3 Loopback Capability” on page 562.

For more information about using operational mode commands to test interfaces, see the *JUNOS Network and Services Interfaces Command Reference*.

Tracing Operations of an Individual Router Interface

To trace the operations of individual routing platform interfaces, include the traceoptions statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
traceoptions {
    flag flag <disable>;
}
```

You can specify the following interface tracing flags:

all—Trace all interface operations.

event—Trace all interface events.

ipc—Trace all interface interprocess communication (IPC) messages.

media—Trace all interface media changes.

The interfaces traceoptions statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system syslog files.

For more information about trace operations, see “Tracing Interface Operations” on page 153.

Damping Interface Transitions

By default, when an interface changes from being up to being down, or from down to up, this transition is advertised immediately to the hardware and the JUNOS software. In some situations—for example, when an interface is connected to an add-drop multiplexer (ADM) or wavelength-division multiplexer (WDM), or to protect against SONET/SDH framer holes—you might want to damp interface transitions. This means not advertising the interface's transition until a certain period of time has passed, called the *hold-time*. When you have damped interface transitions and the interface goes from up to down, the interface is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly when an interface goes from down to up, it is not advertised as being up until it has remained up for the hold-time period.

To damp interface transitions, include the hold-time statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
hold-time up milliseconds down milliseconds;
```

The time can be a value from 0 through 65,534 milliseconds. The default value is 0, which means that interface transitions are not damped. The JUNOS software advertises the transition within 100 milliseconds of the time value you specify.

For most Ethernet interfaces, hold timers are implemented using a one-second polling algorithm. For 1-port, 2-port, and 4-port Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs), hold timers are interrupt driven.

Configuring Multiservice Physical Interface Properties

The adaptive services (AS), collector, monitoring services, and monitoring services II interfaces are multiservice interfaces specifically designed to enable IP services. To configure multiservice physical interface properties on the collector, monitoring services, and AS interfaces, include the multiservice-options statement:

```
multiservice-options {
    boot-command filename;
    (core-dump | no-core-dump);
    (syslog | no-syslog);
}
```

You can include these statements at the following hierarchy levels:

```
[edit interfaces cp-fpc/pic/port]

[edit interfaces mo-fpc/pic/port]

[edit interfaces sp-fpc/pic/port]
```

For more information about the services interfaces, see the *JUNOS Services Interfaces Configuration Guide*.

Enabling or Disabling SNMP Notifications on Physical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the physical interface, include the traps statement at the [edit interfaces *interface-name*] hierarchy level. To disable these notifications on the physical interface, include the no-traps statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
(traps | no-traps);
```

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration. To do this, include the disable statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
disable;
```

Example: Disabling a Physical Interface

Disable a physical interface:

```
[edit interfaces]
so-1/1/0 {
    mtu 8000;
    clocking internal;
    encapsulation ppp;
    sonet-options {
        fcs 16;
    }
    unit 0 {
        family inet {
            address 172.16.0.0/12 {
                destination 172.16.0.4;
            }
        }
    }
}
[edit interfaces]
user@host# set so-1/1/0 disable
[edit interfaces]
user@host# show so-1/1/0
so-1/1/0 {
    disable;       # Interface is marked as disabled
    mtu 8000;
    clocking internal;
    encapsulation ppp;
    sonet-options {
        fcs 16;
    }
```

```
    unit 0 {
        family inet {
            address 172.16.0.0 {
                destination 172.16.0.3;
            }
        }
    }
}
```