



Junos[®] OS

Management Ethernet Interface Feature Guide for Routing Devices



Published: 2013-08-01

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Junos® OS Management Ethernet Interface Feature Guide for Routing Devices

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	vii
	Documentation and Release Notes	vii
	Supported Platforms	vii
	Using the Examples in This Manual	vii
	Merging a Full Example	viii
	Merging a Snippet	viii
	Documentation Conventions	ix
	Documentation Feedback	xi
	Requesting Technical Support	xi
	Self-Help Online Tools and Resources	xi
	Opening a Case with JTAC	xii
Part 1	Overview	
Chapter 1	Management Ethernet Interface	3
	Management Ethernet Interface Overview	3
Part 2	Configuration	
Chapter 2	Management Ethernet Interface	7
	Configuring a Consistent Management IP Address	7
	Configuring the MAC Address on the Management Ethernet Interface	8
Chapter 3	Network Interfaces Configuration Statements and Hierarchy	11
	[edit interfaces] Hierarchy Level	11
	[edit logical-systems] Hierarchy Level	27
Chapter 4	Statement Summary	33
	fast-aps-switch	33
	mac	34
	master-only	34
Part 3	Administration	
Chapter 5	Monitoring Commands	37
	show interfaces (M Series and T Series Routers, and PTX Series Packet Transport Routers Management and Internal Ethernet)	38
Chapter 6	Command Summary	55
	Ethernet Interface Operational Mode Commands	55

Part 4	Troubleshooting	
Chapter 7	Ethernet	63
	tracertoe ethernet	64
Chapter 8	Interface Diagnostics	67
	Interface Diagnostics	67
	Configuring Loopback Testing	67
	Interface Diagnostics	69
	Starting and Stopping a BERT Test	73
	Example: Configuring Bit Error Rate Testing	73
Part 5	Index	
	Index	77

List of Tables

	About the Documentation	vii
	Table 1: Notice Icons	ix
	Table 2: Text and Syntax Conventions	ix
Part 3	Administration	
Chapter 5	Monitoring Commands	37
	Table 3: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields	39
Chapter 6	Command Summary	55
	Table 4: Ethernet Interface Operational Mode Commands	55
Part 4	Troubleshooting	
Chapter 7	Ethernet	63
	Table 5: traceroute ethernet Output Fields	64
Chapter 8	Interface Diagnostics	67
	Table 6: Loopback Modes by Interface Type	68
	Table 7: BERT Capabilities by Interface Type	72

About the Documentation

- Documentation and Release Notes on page vii
- Supported Platforms on page vii
- Using the Examples in This Manual on page vii
- Documentation Conventions on page ix
- Documentation Feedback on page xi
- Requesting Technical Support on page xi

Documentation and Release Notes

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- M Series
- MX Series
- T Series
- J Series

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```


2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page ix defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2 on page ix defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>

- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

- [Management Ethernet Interface on page 3](#)

CHAPTER 1

Management Ethernet Interface

- [Management Ethernet Interface Overview on page 3](#)

Management Ethernet Interface Overview

The router's management Ethernet interface, **fxp0** or **em0**, is an out-of-band management interface that needs to be configured only if you want to connect to the router through the management port on the front of the router. You can configure an IP address and prefix length for this interface, which you commonly do when you first install the Junos OS:

```
[edit]
user@host# set interfaces (fxp0 | em0) unit 0 family inet address/prefix-length
[edit]
user@host# show
interfaces {
  (fxp0 | em0) {
    unit 0 {
      family inet {
        address/prefix-length;
      }
    }
  }
}
```

To determine which management interface type is supported on a router, locate the router and Routing Engine combination in *Supported Routing Engines by Chassis* and note its management Ethernet interface type, either **em0** or **fxp0**.

Related Documentation

- [Configuring a Consistent Management IP Address on page 7](#)
- [Configuring the MAC Address on the Management Ethernet Interface on page 8](#)
- *Configuring MAC Filtering on PTX Series Packet Transport Routers*
- *Ethernet Interfaces*

PART 2

Configuration

- [Management Ethernet Interface on page 7](#)
- [Network Interfaces Configuration Statements and Hierarchy on page 11](#)
- [Statement Summary on page 33](#)

CHAPTER 2

Management Ethernet Interface

- [Configuring a Consistent Management IP Address on page 7](#)
- [Configuring the MAC Address on the Management Ethernet Interface on page 8](#)

Configuring a Consistent Management IP Address

On routers with multiple Routing Engines, each Routing Engine is configured with a separate IP address for the management Ethernet interface. To access the master Routing Engine, you must know which Routing Engine is active and use the appropriate IP address.

Optionally, for consistent access to the master Routing Engine, you can configure an additional IP address and use this address for the management interface regardless of which Routing Engine is active. This additional IP address is active only on the management Ethernet interface for the master Routing Engine. During switchover, the address moves to the new master Routing Engine.



NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is `fxp0`. For TX Matrix Plus routers and T1600 or T4000 routers configured in a routing matrix, the management Ethernet interface is `em0`.



NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the `fxp0` management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the `fxp0` management Ethernet interface so that the commands reference the `em0` management Ethernet interface instead.

To configure an additional IP address for the management Ethernet interface, include the **master-only** statement at the `[edit groups]` hierarchy level.

In the following example, IP address **10.17.40.131** is configured for both Routing Engines and includes a **master-only** statement. With this configuration, the **10.17.40.131** address is active only on the master Routing Engine. The address remains consistent regardless

of which Routing Engine is active. IP address **10.17.40.132** is assigned to **fxp0** on **re0**, and address **10.17.40.133** is assigned to **fxp0** on **re1**.

```
[edit groups re0 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.132/25;
  }
}
[edit groups re1 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.133/25;
  }
}
```

This feature is available on all routers that include dual Routing Engines. On the TX Matrix router, this feature is applicable to the switch-card chassis (SCC) only.

**Related
Documentation**

- [Management Ethernet Interface Overview on page 3](#)
- [Configuring the MAC Address on the Management Ethernet Interface on page 8](#)
- [Ethernet Interfaces](#)

Configuring the MAC Address on the Management Ethernet Interface

By default, the router's management Ethernet interface uses as its MAC address the MAC address that is burned into the Ethernet card.



NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is **fxp0**. For TX Matrix Plus routers and T1600 routers configured in a routing matrix, and TX Matrix Plus routers with 3D SIBs, T1600 routers, and T4000 routers configured in a routing matrix, the management Ethernet interface is **em0**.



NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the **fxp0** management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the **fxp0** management Ethernet interface so that the commands reference the **em0** management Ethernet interface instead.

To display the MAC address used by the router's management Ethernet interface, enter the **show interface fxp0** or **show interface em0** operational mode command.

To change the management Ethernet interface's MAC address, include the **mac** statement at the **[edit interfaces fxp0]** or **[edit interfaces em0]** hierarchy level:

```
[edit interfaces (fxp0 | em0)]
  mac mac-address;
```

Specify the MAC address as six hexadecimal bytes in one of the following formats:
nnnn.nnnn.nnnn (for example, 0011.2233.4455) or *nn:nn:nn:nn:nn:nn* (for example, 00:11:22:33:44:55).



NOTE: If you integrate a standalone T640 router into a routing matrix, the PIC MAC addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, if you integrate a standalone T1600 router into a routing matrix, the PIC MAC addresses for the integrated T1600 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the configuration of a formerly standalone T1600 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

Related Documentation

- [Management Ethernet Interface Overview on page 3](#)
- [Configuring a Consistent Management IP Address on page 7](#)
- [Configuring MAC Filtering on PTX Series Packet Transport Routers](#)
- [Ethernet Interfaces](#)

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

- [\[edit interfaces\] Hierarchy Level on page 11](#)
- [\[edit logical-systems\] Hierarchy Level on page 27](#)

[\[edit interfaces\] Hierarchy Level](#)

The statements at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level can also be configured at the `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]` hierarchy level.



NOTE: The *accounting-profile* statement is an exception to this rule. The *accounting-profile* statement can be configured at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level, but it cannot be configured at the `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]` hierarchy level.

```
interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    account-layer2-overhead (Interface Level) {
      value;
      egress bytes;
      ingress bytes;
    }
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
          (revertive | non-revertive);
        }
      }
    }
  }
}
```

```
    periodic interval;  
    system-priority priority;  
}  
link-protection;  
link-speed speed;  
(loopback | no-loopback);  
mc-ae {  
    chassis-id chassis-id;  
    mc-ae-id mc-ae-id;  
    mode (active-active | active-standby);  
    redundancy-group group-id;  
    status-control (active | standby);  
}  
minimum-links number;  
source-address-filter {  
    mac-address;  
}  
(source-filtering | no-source-filtering);  
}  
shared-scheduler;  
aggregated-sonet-options {  
    link-speed speed | mixed;  
    minimum-links number;  
}  
atm-options {  
    cell-bundle-size cells;  
    ilmi;  
    linear-red-profiles profile-name {  
        high-plp-max-threshold percent;  
        low-plp-max-threshold percent;  
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;  
    }  
}  
mpls {  
    pop-all-labels {  
        required-depth number;  
    }  
}  
pic-type (atm1 | atm2);  
plp-to-clp;  
promiscuous-mode {  
    vpi vpi-identifier;  
}  
scheduler-maps map-name {  
    forwarding-class class-name {  
        epd-threshold cells plp1 cells;  
        linear-red-profile profile-name;  
        priority (high | low);  
        transmit-weight (cells number | percent number);  
    }  
    vc-cos-mode (alternate | strict);  
}  
use-null-cw;  
vpi vpi-identifier {  
    maximum-vcs maximum-vcs;  
    oam-liveness {  
        down-count cells;
```



```

        up-count cells;
    }
    oam-period (seconds | disable);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
         burst length);
        queue-length number;
    }
}
}
clocking clock-source;
data-input (system | interface interface-name);
dce;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
description text;
dialer-options {
    pool pool-name <priority priority>;
}

```

```
disable;
ds0-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback payload;
    start-end-flag (filler | shared);
}
e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (16 | 32);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback (local | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
e3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    framing feet;
    compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
    fcs (16 | 32);
    framing (g.751 | g.832);
    idle-cycle-flag (filler | shared);
    invert-data;
    loopback (local | remote);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
    (unframed | no-unframed);
}
encapsulation type;
es-options {
    backup-interface es-fpc/pic/port;
}
fastether-options {
    802.3ad aex;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
}
source-address-filter {
    mac-address;
```

15

```

frame-length (32 | 64 | 128 | 256);
frame-synchronization {
    alpha number;
    beta number;
    gamma number;
}
minimum-links number;
symmetry (symmetrical-config-and-operation |
    symmetrical-config-asymmetrical-operation);
test-procedure {
    ima-test-start;
    ima-test-stop;
    interface name;
    pattern number;
    period number;
}
transmit-clock (common | independent);
version (1.0 | 1.1);
}
ima-link-options group-id group-id;
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
    interface interface-name {
        (unit unit-number);
    }
}
}
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | nil | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
    lmi-type (ansi | itu | c-lmi);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
}
mac mac-address;

```

```

mlfr-uni-nni-bundle-options {
  acknowledge-retries number;
  acknowledge-timer milliseconds;
  action-red-differential-delay (disable-tx | remove-link);
  drop-timeout milliseconds;
  fragment-threshold bytes;
  cisco-interoperability send-lip-remove-link-for-link-reject;
  hello-timer milliseconds;
  link-layer-overhead percent;
  lmi-type (ansi | itu | c-lmi);
  minimum-links number;
  mrru bytes;
  n391 number;
  n392 number;
  n393 number;
  red-differential-delay milliseconds;
  t391 seconds;
  t392 seconds;
  yellow-differential-delay milliseconds;
}
modem-options {
  dialin (console | routable);
  init-command-string initialization-command-string;
}
mtu bytes;
multi-chassis-protection {
  peer a.b.c.d {
    interface interface-name;
  }
}
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
  interface-type type;
}
no-vpivci-swapping;
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  tx-power dbm;
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
otn-options {
  bytes transmit-payload-type value;
  fec (efec | gfec | gfec-sdfec | none);
  (is-ma | no-is-ma);
  (laser-enable | no-laser-enable);
}

```

```
(line-loopback | no-line-loopback);
(local-loopback | no-local-loopback);
(odu-ttim-action-enable | no-odu-ttim-action-enable);
(otu-ttim-action-enable | no-otu-ttim-action-enable);
odu-delay-management {
    (bypass | no-bypass);
    (monitor-end-point | no-monitor-end-point);
    (number-of-frames | no-number-of-frames);
    (start-measurement | no-start-measurement);
}
(prbs | no-prbs);
preemptive-fast-reroute {
    (backward-frr-enable | no-backward-frr-enable);
    (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
}
rate {
    (fixed-stuff-bytes | no-fixed-stuff-bytes);
    otu4;
    (pass-through | no-pass-through);
}
signal-degrade {
    ber-threshold-clear value;
    ber-threshold-signal-degrade value;
    interval value;
}
trigger trigger-identifier;
tti tti-identifier;
}
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;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    no-termination-request;
    pap {
        access-profile name;
        local-name name;
        local-password password;
        compression;
    }
}
psn-vcip psn-vci-identifier;
psn-vpip psn-vpi-identifier;
receive-bucket {
    overflow (discard | tag);
```

```

    rate percentage;
    threshold bytes;
}
redundancy-options {
    priority sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby;
}
satop-options {
    payload-size n;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
services-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    session-limit {
        maximum number;
    }
}

```

```
    rate new-sessions-per-second;
  }
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-number;
      services priority-level;
    }
  }
}
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote);
  snr-margin {
    current margin;
    snext margin;
  }
}
sonet-options {
  aggregate asx;
  aps {
    advertise-interval milliseconds;
    annex-b;
    authentication-key key;
    fast-aps-switch;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    preserve-interface;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
  }
  bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
  }
  fcs (16 | 32);
  loopback (local | remote);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
}
path-trace trace-string;
(payload-scrambler | no-payload-scrambler);
```



```

    rfc-2615;
    trigger {
        defect ignore;
        hold-time up milliseconds down milliseconds;
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
}
speed (10m | 100m | 1g | oc3 | oc12 | oc48);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        speed (10m | 100m | 1g);
        link-mode (full-duplex | half-duplex);
    }
}
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flag (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
t3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
        value>;
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}

```

```
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address {
            policer {
                input cos-policer-name;
                output cos-policer-name;
            }
        }
    }
}
account-layer2-overhead {
    value;
    egress bytes;
    ingress bytes;
}
accounting-profile name;
advisory-options {
    downstream-rate rate;
    upstream-rate rate;
}
allow-any-vci;
atm-scheduler-map (map-name | default);
backup-options {
    interface interface-name;
}
bandwidth rate;
cell-bundle-size cells;
clear-dont-fragment-bit;
compression {
    rtp {
        f-max-period number;
        maximum-contexts number <force>;
        queues [ queue-numbers ];
        port {
            minimum port-number;
            maximum port-number;
        }
    }
}
compression-device interface-name;
copy-tos-to-outer-ip-header;
demux-destination family;
demux-source family;
demux-options {
    underlying-interface interface-name;
}
description text;
```

```

interface {
    l2tp-interface-id name;
    (dedicated | shared);
}
dialer-options {
    activation-delay seconds;
    callback;
    callback-wait-period time;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    incoming-map {
        caller (caller-id | accept-all);
        initial-route-check seconds;
        load-interval seconds;
        load-threshold percent;
        pool pool-name;
        redial-delay time;
        watch-list {
            [ routes ];
        }
    }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;

```

```

multipoint;
oam-liveness {
    down-count cells;
    up-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
        pap;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
    dynamic-profile profile-name;
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
    burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;

```

```

(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id];
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags-outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id];
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
access-concentrator name;
address address {
    destination address;
}
bundle ml-fpc/pic/port | ls-fpc/pic/port);
duplicate-protection;
dynamic-profile profile-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [filter-names ];
        output filter-name;
    }
    output-list {
        [filter-names ];
    }
}
}
ipsec-sa sa-name;
keep-address-and-control;
max-sessions number;
max-sessions-vs-a-ignore;
mtu bytes;
multicast-only;
negotiate-address;
no-redirects;
policer {
    arp policer-template-name;

```

```

    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-names <service-filter filter-name>;
    }
}
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name <destination address destination-profile
    profile-name | preferred-source-address address>;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
    }
}

```

```

        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                bandwidth-threshold bits-per-second {
                    priority;
                }
            }
            interface priority;
        }
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
}
}
}
}
}
}
}

```

- Related Documentation**
- [Junos OS Hierarchy and RFC Reference](#)
 - [Ethernet Interfaces](#)
 - [Junos OS Network Interfaces Library for Routing Devices](#)

[\[edit logical-systems\] Hierarchy Level](#)

The following lists the statements that can be configured at the **[edit logical-systems]** hierarchy level that are also documented in this manual. For more information about logical systems, see the *Logical Systems Feature Guide for Routing Devices*.

```

logical-systems logical-system-name {
    interfaces interface-name {
        unit logical-unit-number {
            accept-source-mac {
                mac-address mac-address {
                    policer {
                        input cos-policer-name;
                        output cos-policer-name;
                    }
                }
            }
        }
    }
}

```

```
allow-any-vci;
atm-scheduler-map (map-name | default);
bandwidth rate;
backup-options {
    interface interface-name;
}
cell-bundle-size cells;
clear-dont-fragment-bit;
compression {
    rtp {
        f-max-period number;
        port {
            minimum port-number;
            maximum port-number;
        }
        queues [ queue-numbers ];
    }
}
compression-device interface-name;
description text;
interface {
    l2tp-interface-id name;
    (dedicated | shared);
}
dialer-options {
    activation-delay seconds;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    initial-route-check seconds;
    load-threshold number;
    pool pool;
    remote-name remote-callers;
    watch-list {
        [ routes ];
    }
}
disable;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
input-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
```



```

layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
}
dynamic-profile profile-name;
pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
}
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
    burst length);
    queue-length number;
}

```

```
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id]
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id]
vpi vpi-identifier;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
bundle interface-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [filter-names ];
    }
    output filter-name;
    output-list {
        [filter-names ];
    }
}
ipsec-sa sa-name;
keep-address-and-control;
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name> {
```

```

    <mode loose>;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                interface priority;
            }
        }
    }
}

```

```
        bandwidth-threshold bits-per-second {  
            priority;  
        }  
    }  
    route ip-address/mask routing-instance instance-name priority-cost cost;  
    }  
    }  
    virtual-address [ addresses ];  
    }  
    }  
    }  
    }
```

**Related
Documentation**

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

CHAPTER 4

Statement Summary

fast-aps-switch

Syntax	fast-aps-switch;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	(M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only and EX Series switches) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.



NOTE:

- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
 - When the fast-aps-switch statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
 - To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
 - The fast-aps-switch statement cannot be configured when the APS annex-b option is configured.
 - The interfaces that have the fast-aps-switch statement configured cannot be used in virtual private LAN service (VPLS) environments.
-

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Reducing APS Switchover Time in Layer 2 Circuits</i>

mac

Syntax	<code>mac mac-address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the MAC address of the interface. You can configure the MAC address on the management Ethernet interface (fxp0 or em0) only.
Options	mac-address —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55 .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the MAC Address on the Management Ethernet Interface on page 8

master-only

Syntax	<code>master-only;</code>
Hierarchy Level	[edit groups rex interfaces (fxp0 em0) unit <i>logical-unit-number</i> family <i>family</i> address], [edit groups rex logical-systems <i>logical-system-name</i> interfaces fxp0 unit <i>logical-unit-number</i> family <i>family</i> address], [edit interfaces (fxp0 em0) unit <i>logical-unit-number</i> family <i>family</i> address], [edit logical-systems <i>logical-system-name</i> interfaces fxp0 unit <i>logical-unit-number</i> family <i>family</i> address]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the IP address to be used when the Routing Engine is the current master.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Consistent Management IP Address on page 7• <i>CLI User Guide</i>

PART 3

Administration

- [Monitoring Commands on page 37](#)
- [Command Summary on page 55](#)

CHAPTER 5

Monitoring Commands

show interfaces (M Series and T Series Routers, and PTX Series Packet Transport Routers Management and Internal Ethernet)

Syntax	<code>show interfaces <i>interface-name</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Syntax (PTX Series Packet Transport Routers)	<code>show interfaces <i>interface-name</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><statistics></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.
Description	(M Series, T Series, TX Matrix Plus, and PTX Series devices only) Display status information about the management Ethernet and internal Ethernet interfaces.
Options	<p>interface-name—Specify one of the following management interface names: fxp0, fxp1, fxp2, ixgbe0, ixgbe1, bcm0, em0, or em1. For supported Ethernet interface by chassis and Routing Engine, see <i>Supported Routing Engines by Chassis</i>.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces brief (Management Ethernet) on page 42</p> <p>show interfaces (Management Ethernet) on page 42</p> <p>show interfaces (Management Ethernet [TX Matrix Plus Router]) on page 42</p> <p>show interfaces (Management Ethernet [PTX Series Packet Transport Routers]) on page 43</p> <p>show interfaces detail (Management Ethernet) on page 43</p> <p>show interfaces detail (Management Ethernet [TX Matrix Plus Router]) on page 44</p> <p>show interfaces detail (Management Ethernet [PTX Packet Transport Routers]) on page 44</p> <p>show interfaces extensive (Management Ethernet) on page 45</p> <p>show interfaces extensive (Management Ethernet [TX Matrix Plus Router]) on page 46</p>

[show interfaces extensive \(Management Ethernet \[PTX Series Packet Transport Routers\]\) on page 47](#)
[show interfaces brief \(Management Ethernet\) on page 47](#)
[show interfaces brief \(Management Ethernet \[TX Matrix Plus Router\]\) on page 48](#)
[show interfaces brief \(Management Ethernet \[PTX Series Packet Transport Routers\]\) on page 48](#)
[show interfaces \(Internal Ethernet\) on page 48](#)
[show interfaces \(Internal Ethernet \[TX Matrix Plus Router\]\) on page 49](#)
[show interfaces detail \(Internal Ethernet\) on page 49](#)
[show interfaces detail \(Internal Ethernet \[TX Matrix Plus Router\]\) on page 50](#)
[show interfaces extensive \(internal Ethernet\) on page 51](#)
[show interfaces extensive \(internal Ethernet \[TX Matrix Plus Router\]\) on page 52](#)

Output Fields Table 3 on page 39 lists the output fields for the **show interfaces** (management) command on the M Series routers, T Series routers, TX Matrix Plus routers, and PTX Series Packet Transport Routers. Output fields are listed in the approximate order in which they appear.

Table 3: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU)—Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link type	Data transmission type.	detail extensive none

Table 3: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Physical info	Information about the physical interface.	detail extensive
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Media access control (MAC) address of the interface.	detail extensive none
Alternate link address	Backup link address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input packets	Number of packets received on the physical interface.	None specified
Output packets	Number of packets transmitted on the physical interface.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the logical and physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • Resource errors—Sum of transmit drops. 	extensive

Table 3: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface; values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Encapsulation	Encapsulation on the logical interface.	detail extensive none
inet	IP address of the logical interface.	brief
Protocol	Protocol family configured on the logical interface (such as iso or inet6).	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Route table in which this address exists. For example, Route table:0 refers to inet.0.	detail extensive
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none

Table 3: M Series and T Series Router Management and Internal Ethernet show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Broadcast	Broadcast address.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces brief (Management Ethernet)

```

user@host> show interfaces fxp0 brief
Physical interface: fxp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

Logical interface fxp0.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet 192.168.70.143/21

```

show interfaces (Management Ethernet)

```

user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Current address: 00:a0:a5:56:01:89, Hardware address: 00:a0:a5:56:01:89
  Last flapped   : Never
  Input packets  : 80804
  Output packets : 1105

Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500
  Flags: Is-Primary
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.168.64/21, Local: 192.168.70.143,
    Broadcast: 192.168.71.255

```

show interfaces (Management Ethernet [TX Matrix Plus Router])

```

user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:80:f9:26:00:c0, Hardware address: 00:80:f9:26:00:c0
  Last flapped   : Never
  Input packets  : 1424
  Output packets : 5282

```

```

Logical interface em0.0 (Index 3) (SNMP ifIndex 18)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 1424
  Output packets: 5282
  Protocol inet, MTU: 1500
  Flags: Is-Primary
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
    192.168.178.127

```

show interfaces (Management Ethernet [PTX Series Packet Transport Routers])

```

user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
  Last flapped   : Never
  Input packets  : 212581
  Output packets : 71

    Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 212551
    Output packets: 71
    Protocol inet, MTU: 1500
    Flags: Is-Primary
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 192.168.3/24, Local: 192.168.3.30,
      Broadcast: 192.168.3.255

```

show interfaces detail (Management Ethernet)

```

user@host> show interfaces fxp0 detail
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info   : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:a0:a5:56:01:89, Hardware address: 00:a0:a5:56:01:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 6484031
    Output bytes : 167503
    Input packets: 81008
    Output packets: 1110

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
    Flags: Is-Primary

```

```
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.168.64/21, Local: 192.168.70.143,
Broadcast: 192.168.71.255, Generation: 1
```

show interfaces detail (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 detail
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:80:f9:26:00:c0, Hardware address: 00:80:f9:26:00:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :           124351
    Output bytes  :          1353212
    Input packets :           1804
    Output packets:           5344
  IPv6 transit statistics:
    Input bytes   :           0
    Output bytes  :           0
    Input packets :           0
    Output packets:           0

Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
    Input bytes   :           117135
    Output bytes  :          1331647
    Input packets :           1804
    Output packets:           5344
  Local statistics:
    Input bytes   :           117135
    Output bytes  :          1331647
    Input packets :           1804
    Output packets:           5344
  Protocol inet, MTU: 1500, Generation: 1, Route table: 0
  Flags: Is-Primary
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
    192.168.178.127, Generation: 1
```

show interfaces detail (Management Ethernet [PTX Packet Transport Routers])

```
user@host> show interfaces detail em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 0, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
```



```

Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
Alternate link address: Unspecified
Last flapped    : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :          15255909
  Output bytes  :           4608
  Input packets :          214753
  Output packets:           72
IPv6 transit statistics:
  Input bytes   :           0
  Output bytes  :           0
  Input packets :           0
  Output packets:           0

Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes   :          14394630
  Output bytes  :           3024
  Input packets :          214723
  Output packets:           72
Local statistics:
  Input bytes   :          14394630
  Output bytes  :           3024
  Input packets :          214723
  Output packets:           72
Protocol inet, MTU: 1500, Generation: 1, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 192.168.3/24, Local: 192.168.3.30,
  Broadcast: 192.168.3.255, Generation: 1

```

show interfaces extensive (Management Ethernet)

```

user@host> show interfaces fxp0 extensive
Physical interface: fxp0, Enabled, Physical link is Up
Interface index: 1, SNMP ifIndex: 1, Generation: 0
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
Device flags      : Present Running
Interface flags: SNMP-Traps
Link type         : Half-Duplex
Physical info      : Unspecified
Hold-times        : Up 0 ms, Down 0 ms
Current address: 00:a0:a5:56:01:89, Hardware address: 00:a0:a5:56:01:89
Alternate link address: Unspecified
Last flapped      : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :          6678904
  Output bytes  :          169657
  Input packets :          83946
  Output packets:          1127
Input errors:
  Errors: 12, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

```

```
Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 6, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.168.64/21, Local: 192.168.70.143,
Broadcast: 192.168.71.255, Generation: 1
```

show interfaces extensive (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 extensive
```

```
Physical interface: em0, Enabled, Physical link is Up
Interface index: 8, SNMP ifIndex: 17, Generation: 2
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
Device flags      : Present Running
Interface flags: SNMP-Traps
Link type        : Full-Duplex
Physical info     : Unspecified
Hold-times       : Up 0 ms, Down 0 ms
Current address: 00:80:f9:26:00:c0, Hardware address: 00:80:f9:26:00:c0
Alternate link address: Unspecified
Last flapped     : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes      :          127120
Output bytes     :        1357414
Input packets    :           1843
Output packets   :          5372
IPv6 transit statistics:
Input bytes      :              0
Output bytes     :              0
Input packets    :              0
Output packets   :              0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
```

```
Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes      :          119748
Output bytes     :        1335719
Input packets    :           1843
Output packets   :          5372
Local statistics:
Input bytes      :          119748
Output bytes     :        1335719
Input packets    :           1843
Output packets   :          5372
Protocol inet, MTU: 1500, Generation: 1, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1
```

show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers])

```

user@host> show interfaces extensive em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 0, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:80:f9:25:00:1b, Hardware address: 00:80:f9:25:00:1b
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          15236459
    Output bytes  :           4608
    Input packets :         214482
    Output packets:           72
  IPv6 transit statistics:
    Input bytes   :           0
    Output bytes  :           0
    Input packets :           0
    Output packets:           0
  Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
  Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
    Input bytes   :         14376264
    Output bytes  :          3024
    Input packets :         214452
    Output packets:           72
  Local statistics:
    Input bytes   :         14376264
    Output bytes  :          3024
    Input packets :         214452
    Output packets:           72
  Protocol inet, MTU: 1500, Generation: 1, Route table: 0
  Flags: Is-Primary
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
    Destination: 192.168.3/24, Local: 192.168.3.30,
    Broadcast: 192.168.3.255, Generation: 1

```

show interfaces brief (Management Ethernet)

```

user@host> show interfaces fxp1 brief
Physical interface: fxp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

```

```
Logical interface fxp1.0
Flags: SNMP-Traps Encapsulation: ENET2
inet 10.0.0.4/8
inet6 fe80::200:ff:fe00:4/64
fec0::10:0:0:4/64
tnp 4
```

show interfaces brief (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps

Logical interface em0.0
Flags: SNMP-Traps Encapsulation: ENET2
inet 192.168.178.11/25
```

show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,

Speed: 1000mbps
Device flags : Present Running
Interface flags: SNMP-Traps

Logical interface em0.0
Flags: SNMP-Traps Encapsulation: ENET2
inet 192.168.3.30/24
```

```
root@aboslutely> show interfaces em0 terse
Interface      Admin Link Proto  Local      Remote
em0            up    up
em0.0          up    up  inet    192.168.3.30/24
```

show interfaces (Internal Ethernet)

```
user@host> show interfaces fxp1
Physical interface: fxp1, Enabled, Physical link is Up
Interface index: 2, SNMP ifIndex: 2
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Current address: 02:00:00:00:00:04, Hardware address: 02:00:00:00:00:04
Last flapped : Never
Input packets : 30655
Output packets: 33323

Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500
Flags: Is-Primary
Addresses, Flags: Is-Default Is-Preferred Is-Primary
Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255
Protocol inet6, MTU: 1500
Flags: Is-Primary
Addresses, Flags: Is-Preferred
```

```

        Destination: fe80::/64, Local: fe80::200:ff:fe00:4
Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4
Protocol tnp, MTU: 1500
Flags: Primary, Is-Primary
Addresses
Local: 4

```

show interfaces (Internal Ethernet [TX Matrix Plus Router])

```

user@host> show interfaces ixgbe0
Physical interface: ixgbe0, Enabled, Physical link is Up
Interface index: 2, SNMP ifIndex: 116
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
Link type      : Full-Duplex
Current address: 02:00:00:22:00:04, Hardware address: 02:00:00:22:00:04
Last flapped   : Never
Input packets  : 2301738
Output packets : 3951155

Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117)
Flags: SNMP-Traps Encapsulation: ENET2
Input packets : 2301595
Output packets: 3951155
Protocol inet, MTU: 1500
Flags: Is-Primary
Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255
Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 128/2, Local: 162.0.0.4, Broadcast: 191.255.255.255
Protocol inet6, MTU: 1500
Flags: Is-Primary
Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
Protocol tnp, MTU: 1500
Flags: Primary, Is-Primary
Addresses
Local: 0x22000004

```

show interfaces detail (Internal Ethernet)

```

user@host> show interfaces fxp1 detail
Physical interface: fxp1, Enabled, Physical link is Up
Interface index: 2, SNMP ifIndex: 2, Generation: 1
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
Link type      : Full-Duplex
Physical info   : Unspecified
Hold-times     : Up 0 ms, Down 0 ms
Current address: 02:00:00:00:00:04, Hardware address: 02:00:00:00:00:04
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes    :                2339969

```

```

Output bytes :          15880707
Input packets:          30758
Output packets:         33443

Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 7, Route table: 1
  Flags: Is-Primary
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
    Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
    Generation: 3
  Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
      Broadcast: Unspecified, Generation: 5
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
      Generation: 7
  Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
    Flags: Primary, Is-Primary
    Addresses, Flags: None
      Destination: Unspecified, Local: 4, Broadcast: Unspecified,
      Generation: 8

```

show interfaces detail (Internal Ethernet [TX Matrix Plus Router])

```

user@host> show interfaces ixgbe0 detail
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:22:00:04, Hardware address: 02:00:00:22:00:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          238172825
    Output bytes :         1338948955
    Input packets:         2360984
    Output packets:        4061512
  IPv6 transit statistics:
    Input bytes :          0
    Output bytes :          0
    Input packets:          0
    Output packets:          0

Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes :          228720309
  Output bytes :         1261387447
  Input packets:         2360841
  Output packets:        4061512
  IPv6 transit statistics:
    Input bytes :          0

```

```

        Output bytes :           0
        Input  packets:           0
        Output packets:           0
Local statistics:
    Input bytes :       228720309
    Output bytes :       1261387447
    Input  packets:       2360841
    Output packets:       4061512
Protocol inet, MTU: 1500, Generation: 2, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation:
2
        Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 128/2, Local: 162.0.0.4, Broadcast: 191.255.255.255,
Generation: 3
    Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
Generation: 4
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500, Generation: 5
    Generation: 4, Route table: 1
    Flags: Primary, Is-Primary
    Addresses, Flags: None
        Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6

```

show interfaces extensive (internal Ethernet)

```

user@host> show interfaces fxp1 extensive
Physical interface: fxp1, Enabled, Physical link is Up
    Interface index: 2, SNMP ifIndex: 2, Generation: 1
    Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
    Speed: 100mbps
    Device flags : Present Running
    Interface flags: SNMP-Traps
    Link type : Full-Duplex
    Physical info : Unspecified
    Hold-times : Up 0 ms, Down 0 ms
    Current address: 02:00:00:00:00:04, Hardware address: 02:00:00:00:00:04
    Alternate link address: Unspecified
    Last flapped : Never
    Statistics last cleared: Never
    Traffic statistics:
        Input bytes :       2349897
        Output bytes :       15888605
        Input  packets:       30896
        Output packets:       33607
    Input errors:
        Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
        Policed discards: 0, Resource errors: 0
    Output errors:
        Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
        Resource errors: 0

    Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
        Flags: SNMP-Traps Encapsulation: ENET2
        Protocol inet, MTU: 1500, Generation: 7, Route table: 1

```

```

Flags: Is-Primary
Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
  Generation: 3
Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
Flags: Is-Primary
Addresses, Flags: Is-Preferred
  Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
  Broadcast: Unspecified, Generation: 5
Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
  Generation: 7
Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
Flags: Primary, Is-Primary
Addresses, Flags: None
  Destination: Unspecified, Local: 4, Broadcast: Unspecified,
  Generation: 8

```

show interfaces extensive (internal Ethernet [TX Matrix Plus Router])

```

user@host> show interfaces ixgbe0 extensive
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:22:00:04, Hardware address: 02:00:00:22:00:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          242730780
    Output bytes  :          1348312269
    Input packets :          2398737
    Output packets:          4133510
  IPv6 transit statistics:
    Input bytes   :          0
    Output bytes  :          0
    Input packets :          0
    Output packets:          0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes   :          233127252
  Output bytes  :          1269350897
  Input packets :          2398594
  Output packets:          4133510
  IPv6 transit statistics:
    Input bytes   :          0
    Output bytes  :          0

```



```

    Input packets:          0
    Output packets:         0
Local statistics:
  Input bytes :            233127252
  Output bytes :          1269350897
  Input packets:           2398594
  Output packets:          4133510
Protocol inet, MTU: 1500, Generation: 2, Route table: 1
  Flags: Is-Primary
  Addresses, Flags: Is-Preferred
    Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation:
2
    Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
    Destination: 128/2, Local: 162.0.0.4, Broadcast: 191.255.255.255,
Generation: 3
  Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::200:ff:fe22:4
Generation: 4
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
    Destination: fec0::/64, Local: fec0::a:22:0:4
  Protocol tnp, MTU: 1500, Generation: 5
  Generation: 4, Route table: 1
    Flags: Primary, Is-Primary
    Addresses, Flags: None
    Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6

```


CHAPTER 6

Command Summary

- [Ethernet Interface Operational Mode Commands on page 55](#)

Ethernet Interface Operational Mode Commands

[Table 4 on page 55](#) summarizes the command-line interface (CLI) commands that you can use to monitor and troubleshoot aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. Commands are listed in alphabetical order.

Table 4: Ethernet Interface Operational Mode Commands

Task	Command
Clear dynamic VLAN interfaces.	<i>clear auto-configuration interfaces</i>
Clear a specified dynamic agent circuit identifier (ACI) interface set configured on the router. You can clear only those ACI interface sets that have no subscriber interface members.	<i>clear auto-configuration interfaces interface-set</i>
Clear Link Aggregation Control Protocol (LACP) statistics.	<i>clear lacp statistics</i>
Clear Link Aggregation Control Protocol (LACP) timeout entries.	<i>clear lacp timeouts</i>
Clear learned MAC addresses from the hardware and MAC database. Static MAC addresses are not cleared.	<i>clear interfaces mac-database</i>
Clear statistics that are collected for every MAC address, including policer statistics, on a given physical or logical interface.	<i>clear interfaces mac-database statistics</i>
Clear statistics that are collected for interface sets.	<i>clear interfaces interface-set statistics</i>
Clear the existing continuity measurement and restart counting the operational uptime.	<i>clear oam ethernet connectivity-fault-management continuity-measurement</i>

Table 4: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Clear ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) delay statistics and ETH-DM frame counts. (MX Series routers)	<i>clear oam ethernet connectivity-fault-management delay-statistics</i>
Clear Operation, Administration, and Management (OAM) and connectivity fault management (CFM) linktrace database information.	<i>clear oam ethernet connectivity-fault-management linktrace path-database</i>
Clear all loss statistics maintained by CFM for a given maintenance domain and maintenance association.	<i>clear oam ethernet connectivity-fault-management loss-statistics</i>
Clear connectivity-fault-management policer statistics.	<i>clear oam ethernet connectivity-fault-management policer</i>
Clear the Ethernet OAM service-level agreement (SLA) iterator statistics.	<i>clear oam ethernet connectivity-fault-management sla-iterator-statistics</i>
Clear all statistics maintained by CFM. (Routers that support IEEE 802.1ag OAM CFM) In addition, for interfaces that support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), also clear any ETH-DM statistics and frame counts for CFM maintenance association end points (MEPs).	<i>clear oam ethernet connectivity-fault-management statistics</i>
Clear ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) delay statistics and ETH-SLM frame counts. (MX Series routers, Modular Port Concentrators only)	<i>clear oam ethernet connectivity-fault-management synthetic-loss-measurement</i>
Clear Operation, Administration, and Management (OAM) link fault management state information and restart the link discovery process on Ethernet interfaces.	<i>clear oam ethernet link-fault-management state</i>
Clear Operation, Administration, and Management (OAM) statistics link fault management statistics for Ethernet interfaces.	<i>clear oam ethernet link-fault-management statistics</i>
Clear the statistics for all Ethernet ring protection groups or a specific Ethernet ring protection group.	<i>clear protection-group ethernet-ring statistics</i>
Check the reachability of a remote IEEE 802.1ag OAM maintenance association end point (MEP) or maintenance association intermediate point (MIP).	<i>ping ethernet</i>

Table 4: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Manually rebalance the subscribers on an aggregated Ethernet bundle with targeted distribution enabled.	<i>request interface rebalance (Aggregated Ethernet for Subscriber Management)</i>
Manually revert egress traffic from the designated backup link to the designated primary link of an aggregated Ethernet interface for which link protection is enabled, or manually switch egress traffic from the primary link to the backup link.	<i>request interface (revert switchover) (Aggregated Ethernet Link Protection)</i>
Force LACP link switchover.	<i>request lacp link-switchover</i>
Clear the lockout, force switch, manual switch, exercise, and wait-to-restore states.	<i>request protection-group ethernet-aps clear</i>
Test if APS is operating correctly.	<i>request protection-group ethernet-aps exercise</i>
Force traffic to switch from the active path to the alternate path.	<i>request protection-group ethernet-aps force-switch</i>
Lock the protection path, forcing the use of the working path.	<i>request protection-group ethernet-aps lockout</i>
Force traffic to switch from the active path to the alternate path.	<i>request protection-group ethernet-aps manual-switch</i>
Display status information about aggregated Fast Ethernet or Gigabit Ethernet router interfaces.	<i>show interfaces (Aggregated Ethernet)</i> <i>show interfaces (far-end-interval)</i>
Display status information about Fast Ethernet interfaces.	<i>show interfaces (Fast Ethernet)</i>
Display status information about the specified Gigabit Ethernet interface.	<i>show interfaces (Gigabit Ethernet)</i>
Display status information about 10-Gigabit Ethernet router interfaces.	<i>show interfaces (10-Gigabit Ethernet)</i>

Table 4: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display IPv6 interface statistics for IPv6 traffic traversing through the IQ2 and IQ2E PICs on standalone T640 routers and on T640 routers in a TX Matrix or in a TXP Matrix.	<i>show interfaces extensive</i>
Display IPv6 interface statistics for IPv6 traffic traversing through the IQ2 PICs on M10i and M120 routers.	
Display IPv6 interface statistics for IPv6 traffic traversing through the IQ2E PICs on M10i, M120, and M320 routers.	
Display information about Gigabit Ethernet or 10-Gigabit Ethernet router interface sets.	<i>show interfaces interface-set (Ethernet Interface Set)</i>
Display information about Gigabit Ethernet or 10-Gigabit Ethernet router interface set queues.	<i>show interfaces interface-set queue</i>
Display the transceiver temperature, laser bias current, laser output power, receive optical power, and related alarms for 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces.	<i>show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)</i>
Display information about integrated routing and bridging interfaces.	<i>show interfaces irb</i>
Display status information about the distribution of subscribers on different links in an aggregated Ethernet bundle.	<i>show interfaces targeting (Aggregated Ethernet for Subscriber Management)</i>
Display Link Aggregation Control Protocol (LACP) information for aggregated, Fast Ethernet, or Gigabit Ethernet router interfaces.	<i>show lacp interfaces</i>
Display Link Aggregation Control Protocol (LACP) statistics.	<i>show lacp statistics</i>
Display Link Aggregation Control Protocol timeout entries.	<i>show lacp timeouts</i>
Display MAC address information for Gigabit Ethernet router interfaces.	<i>show interfaces mac-database (Gigabit Ethernet)</i>
Display information on a specified interface that is part of a multichassis link aggregation configuration.	<i>show interfaces mc-ae</i>
Display ETH-DM statistics for CFM MEPs. (MX Series routers, Ethernet DPCs).	<i>show oam ethernet connectivity-fault-management delay-statistics</i>

Table 4: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display IEEE 802.1ag OAM connectivity fault management forwarding state information for Ethernet interfaces.	<i>show oam ethernet connectivity-fault-management forwarding-state</i>
Display OAM connectivity fault management information for Ethernet interfaces. For interfaces that support ETH-DM, also display any ETH-DM frame counts when the detail or extensive option is included. In all other cases, ETH-DM frame counts are zero.	<i>show oam ethernet connectivity-fault-management interfaces</i>
Display OAM connectivity fault management linktrace path database information.	<i>show oam ethernet connectivity-fault-management linktrace path-database</i>
Display OAM connectivity fault management maintenance association end point (MEP) database information. For interfaces that support ETH-DM, also display any ETH-DM frame counts. In all other cases, ETH-DM frame counts are zero.	<i>show oam ethernet connectivity-fault-management mep-database</i>
Display ETH-DM statistics and frame counts for CFM MEPs. (MX Series routers, Ethernet DPCs)	<i>show oam ethernet connectivity-fault-management mep-statistics</i>
Display ETH-LM statistics for on-demand mode only.	<i>show oam ethernet connectivity-fault-management loss-statistics</i>
Display information about maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag standard for connectivity fault management (CFM).	<i>show oam ethernet connectivity-fault-management mip</i>
Display OAM connectivity fault management path database information for hosts configured with MEP.	<i>show oam ethernet connectivity-fault-management path-database</i>
Displays connectivity-fault-management policer statistics.	<i>show oam ethernet connectivity-fault-management policer</i>
Display the Ethernet OAM service-level agreement (SLA) iterator statistics.	<i>show oam ethernet connectivity-fault-management sla-iterator-statistics</i>
Display ETH-SLM statistics for CFM MEPs (on-demand mode only). (MX Series routers, Ethernet MPCs)."	<i>show oam ethernet connectivity-fault-management synthetic-loss-statistics</i>

Table 4: Ethernet Interface Operational Mode Commands (*continued*)

Task	Command
Display OAM Ethernet Virtual Connection (EVC) information for hosts configured with Ethernet Local Management Interface (E-LMI). (MX series only)	<i>show oam ethernet evc</i>
Display OAM fault management statistics for Ethernet interfaces.	<i>show oam ethernet link-fault-management</i>
Display OAM Ethernet Local Management Interface status information for an LMI configured interface. (MX series only)	<i>show oam ethernet lmi</i>
Display OAM Ethernet Local Management Interface statistics for an LMI configured interface. (MX series only)	<i>show oam ethernet lmi statistics</i>
Display protection group Ethernet ring Automatic Protection Switching (APS).	<i>show protection-group ethernet-ring aps</i>
Display data channel information for all Ethernet ring protection groups or for a specific Ethernet ring protection group.	<i>show protection-group ethernet-ring data-channel</i>
Display protection group Ethernet ring interfaces.	<i>show protection-group ethernet-ring interface</i>
Display protection group Ethernet ring nodes.	<i>show protection-group ethernet-ring node-state</i>
Display protection group Ethernet ring statistics.	<i>show protection-group ethernet-ring statistics</i>
Display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.	<i>show protection-group ethernet-ring vlan</i>
Trace the path between two Ethernet OAM end points.	<i>traceroute ethernet</i>

PART 4

Troubleshooting

- [Ethernet on page 63](#)
- [Interface Diagnostics on page 67](#)

CHAPTER 7

Ethernet

traceroute ethernet

Syntax	traceroute ethernet (<i>mac-address</i> <i>mep-id</i>) maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> ttl <i>value</i> <wait seconds>
Release Information	Command introduced in Junos OS Release 9.0. mep-id option introduced in Junos OS Release 9.1.
Description	<p>Triggers the linktrace protocol to trace the route between two maintenance points. The result of the traceroute protocol is stored in the path database. To display the path database, use the show oam ethernet connectivity-fault-management path-database command.</p> <p>Before using the traceroute command, you can verify the remote MEP's MAC address using the show oam ethernet connectivity-fault-management path-database command.</p>
Options	<p>mac-address—Destination unicast MAC address of the remote maintenance point.</p> <p>mep-id—MEP identifier of the remote maintenance point. The range of values is 1 through 8191.</p> <p>maintenance-association <i>ma-name</i>—Specifies an existing maintenance association from the set of configured maintenance associations.</p> <p>maintenance-domain <i>md-name</i>—Specifies an existing maintenance domain from the set of configured maintenance domains.</p> <p>ttl <i>value</i>—Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4.</p> <p>wait <i>seconds</i>—(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.</p>
Required Privilege Level	network
List of Sample Output	traceroute ethernet on page 65
Output Fields	<p>Table 5 on page 64 lists the output fields for the traceroute ethernet command. Output fields are listed in the approximate order in which they appear.</p>

Table 5: traceroute ethernet Output Fields

Field Name	Field Description
Linktrace to	MAC address of the destination maintenance point.
Interface	Local interface used to send the linktrace message (LTM).

Table 5: traceroute ethernet Output Fields (*continued*)

Field Name	Field Description
Maintenance Domain	Maintenance domain specified in the traceroute command.
Level	Maintenance domain level configured.
Maintenance Association	Maintenance association specified in the traceroute command.
Local Mep	The local maintenance end point identifier.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all Maintenance Domains. Use the transaction identifier to match an incoming linktrace response (LTR), with a previously sent LTM.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag maintenance point that is sending the linktrace message.
Next-hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.

Sample Output

traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:90:69:7e:01:ff
```

```
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
```

```
Maintenance Domain: MD1, Level: 7
```

```
Maintenance Association: MA1, Local Mep: 1
```

Hop	TTL	Source MAC address	Next hop MAC address
Transaction Identifier:100001			
1	63	00:00:aa:aa:aa:aa	00:00:bb:bb:bb:bb
2	62	00:00:bb:bb:bb:bb	00:00:cc:cc:cc:cc
3	61	00:00:cc:cc:cc:cc	00:01:02:03:04:05
4	60	00:01:02:03:04:05	00:00:00:00:00:00

CHAPTER 8

Interface Diagnostics

- [Interface Diagnostics on page 67](#)

Interface Diagnostics

You can use two diagnostic tools to test the physical layer connections of interfaces: loopback testing and bit error rate test (BERT) testing. Loopback testing enables you to verify the connectivity of a circuit. BERT testing enables you to identify poor signal quality on a circuit. This section contains the following topics:

- [Configuring Loopback Testing on page 67](#)
- [Interface Diagnostics on page 69](#)

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own data link layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own data link layer packets, you can assume the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

There are several types of loopback testing supported by the Junos OS, as follows:

- DCE local—Loops packets back on the local DCE.
- DCE remote—Loops packets back on the remote DCE.

- **Local**—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- **Payload**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- **Remote**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 6 on page 68 shows the loopback modes supported on the various interface types.

Table 6: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	<i>Configuring Ethernet Loopback Capability</i>
Circuit Emulation E1	Local and remote	<i>Configuring E1 Loopback Capability</i>
Circuit Emulation T1	Local and remote	<i>Configuring T1 Loopback Capability</i>
E1 and E3	Local and remote	<i>Configuring E1 Loopback Capability and Configuring E3 Loopback Capability</i>
NxDSO	Payload	<i>Configuring Channelized E1 IQ and IQE Interfaces, Configuring T1 and NxDSO Interfaces, Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode), Configuring Channelized STM1 IQ and IQE Interfaces, and Configuring Channelized T3 IQ Interfaces</i>
Serial (V.35 and X.21)	Local and remote	<i>Configuring Serial Loopback Capability</i>
Serial (EIA-530)	DCE local, DCE remote, local, and remote	<i>Configuring Serial Loopback Capability</i>
SONET/SDH	Local and remote	<i>Configuring SONET/SDH Loopback Capability</i>

Table 6: Loopback Modes by Interface Type (*continued*)

Interface	Loopback Modes	Usage Guidelines
T1 and T3	Local, payload, and remote	<i>Configuring T1 Loopback Capability</i> and <i>Configuring T3 Loopback Capability</i> <i>See also Configuring the T1 Remote Loopback Response</i>

To configure loopback testing, include the **loopback** statement:

loopback mode;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]
- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

Interface Diagnostics

BERT allows you to troubleshoot problems by checking the quality of links. You can configure any of the following interfaces to execute a 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 IQ, and OC12 IQ interfaces.

A BERT test requires a line loop to be in place on either the transmission devices or the far-end router. The local router generates a known bit pattern and sends it out the transmit path. The received pattern is then verified against the sent pattern. The higher the bit error rate of the received pattern, the worse the noise is on the physical circuit. As you move the position of the line loop increasingly downstream toward the far-end router, you can isolate the troubled portion of the link.

To configure BERT, you must configure the duration of the test, the bit pattern to send on the transmit path, and the error rate to monitor when the inbound pattern is received.

To configure the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream, include the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements, respectively, at the [edit interfaces *interface-name* *interface-type*-options] hierarchy level:

```
[edit interfaces interface-name interface-type-options]
bert-algorithm algorithm;
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. 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 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
...
```

For specific hierarchy information, see the individual interface types.



NOTE: The 4-port E1 PIC supports only the following algorithms:

pseudo-2e11-o152	Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151	Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151	Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151	Pattern is 2^23 (per 0.151 standard)

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```

all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7            Pattern is 2^7 - 1
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set

```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The IQE PICs support only the following algorithms:

```

all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 (511 type) standard)
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 and 0.153 (2047 type)
standards)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151      Pattern is 2^23 - 1 (per 0.151 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set

```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

all-ones-repeating	Repeating one bits
all-zeros-repeating	Repeating zero bits
alternating-double-ones-zeros	Alternating pairs of ones and zeros
alternating-ones-zeros	Alternating ones and zeros
repeating-1-in-4	1 bit in 4 is set
repeating-1-in-8	1 bit in 8 is set
repeating-3-in-24	3 bits in 24 are set
pseudo-2e9-o153	Pattern is $2^9 - 1$ (per 0.153 standard)
pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e20-o153	Pattern is $2^{20} - 1$ (per 0.153 standard)
pseudo-2e23-o151	Pattern is $2^{23} - 1$ (per 0.151 standard)

Table 7 on page 72 shows the BERT capabilities for various interface types.

Table 7: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)		<ul style="list-style-type: none"> Limited algorithms
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)		<ul style="list-style-type: none"> Limited algorithms
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> Single port at a time Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> Single port at a time
Channelized OC12	N/A	Yes (channel 0–11)	<ul style="list-style-type: none"> Single channel at a time Limited algorithms No bit count
Channelized STM1	Yes (channel 0–62)	N/A	<ul style="list-style-type: none"> 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)	<ul style="list-style-type: none"> Multiple ports and channels Limited algorithms for T1 No error insert for T1 No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see *Channelized IQ and IQE Interfaces Properties*.

Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

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

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name* *interface-type*-bert-start** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you wish to terminate the test sooner, issue the **test interface *interface-name* *interface-type*-bert-stop** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the *Junos OS Operational Mode Commands*.



NOTE: To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

Example: Configuring Bit Error Rate Testing

Configure a BERT test on a T3 interface. In this example, the run duration lasts for 120 seconds. The configured error rate is 0, which corresponds to a bit error rate of 10^{-0} (1 error per bit). The configured bit pattern of **all-ones-repeating** means that every bit the interface sends is a set to a value of 1.

```
[edit interfaces]
t3-1/2/0 {
  t3-options {
    bert algorithm all-ones-repeating;
    bert-error-rate 0;
    bert-period 120;
```

```
}  
}
```

PART 5

Index

- [Index on page 77](#)

Index

Symbols

#, comments in configuration statements.....	x
(), in syntax descriptions.....	x
< >, in syntax descriptions.....	x
[], in configuration statements.....	x
{ }, in configuration statements.....	x
(pipe), in syntax descriptions.....	x

B

BERT	
configuring interface diagnostics.....	69
bert-algorithm statement	
usage guidelines.....	69
bert-error-rate statement	
usage guidelines.....	69
bert-period statement	
usage guidelines.....	69
bit error rate test See BERT	
braces, in configuration statements.....	x
brackets	
angle, in syntax descriptions.....	x
square, in configuration statements.....	x

C

comments, in configuration statements.....	x
conventions	
text and syntax.....	ix
curly braces, in configuration statements.....	x
customer support.....	xi
contacting JTAC.....	xi

D

documentation	
comments on.....	xi

E

Ethernet interfaces	
status information, displaying	
internal.....	38
management.....	38

F

fast-aps-switch statement.....	33
font conventions.....	ix

I

interfaces	
configuration statements.....	11
internal Ethernet interface	
status information, displaying.....	38
IP addresses	
management Ethernet interface.....	7

L

linktrace	
tracing.....	64
logical systems	
configuration statements.....	27
loopback testing.....	67

M

MAC addresses	
management Ethernet interface.....	8
mac statement.....	34
usage guidelines.....	8
management Ethernet interface	
IP address.....	7
MAC address.....	8
manuals	
comments on.....	xi
master-only statement.....	34
usage guidelines.....	7

P

parentheses, in syntax descriptions.....	x
--	---

S

show interfaces (M Series and T Series Management and Internal Ethernet)	
command.....	38
support, technical See technical support	
syntax conventions.....	ix

T

technical support	
contacting JTAC.....	xi
traceroute ethernet.....	64
traceroute ethernet command.....	64

