Junos OS

Ethernet Interfaces for Routing Devices

Release

13.2

Published: 2013-08-29

Copyright © 2013, Juniper Networks, Inc.

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 Ethernet Interfaces for Routing Devices
Release 13.2
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

Revision History
July 2013—R1 Junos OS 13.2

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.

Abbreviated Table of Contents

Copyright © 2013, Juniper Networks, Inc.

Table of Contents

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

Part 5 Index

 Copyright © 2013, Juniper Networks, Inc.

List of Figures

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

List of Tables

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

About This Guide

This preface provides the following guidelines for using the *Junos OS Ethernet Interfaces for Routing Devices*:

Junos Documentation and Release Notes

For a list of related Junos documentation, see http://www.juniper.net/techpubs/software/junos/.

If the information in the latest release notes differs from the information in the documentation, follow the *Junos 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/.

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at http://www.juniper.net/books.

Objectives

This guide provides an overview of the network interfaces features of the JUNOS Software and describes how to configure these properties on the routing platform.

NOTE: For additional information about the Junos OS—either corrections to or information that might have been omitted from this guide—see the software release notes at http://www.juniper.net/ .

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Supported Routing Platforms

For the features described in this manual, the JUNOS Software currently supports the following routing platforms:

- J Series
- M Series

Copyright © 2013, Juniper Networks, Inc.

- MX Series
- T Series

Using the Indexes

This reference contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines,* refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

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.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
   ```

Copyright © 2013, Juniper Networks, Inc.

```
          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.xsl; }
   ```

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 xxxiii defines notice icons used in this guide.

Copyright © 2013, Juniper Networks, Inc.

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 xxxiii 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`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	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** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• 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 *metric*>;**

Copyright © 2013, Juniper Networks, Inc.

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

Convention	Description	Examples
| (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** **(*string1* | *string2* | *string3*)**
# (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 [*community-ids*]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		    route default {
		      nexthop *address*;
		      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.	• 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,

Copyright © 2013, Juniper Networks, Inc.

or are covered under warranty, and need postsales 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/
- 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, visit us at http://www.juniper.net/support/requesting-support.html

 Copyright © 2013, Juniper Networks, Inc.

PART 1

Ethernet Interfaces Configuration Statements Overview

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 1

Ethernet Interfaces Configuration Statements and Hierarchy

The following interfaces hierarchy listings show the complete configuration statement hierarchy for the indicated hierarchy levels, listing all possible configuration statements within the indicated hierarchy levels, and showing their level in the configuration hierarchy. When you are configuring the Junos OS, your current hierarchy level is shown in the banner on the line preceding the **user@host#** prompt.

This section contains the following topics:

[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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
  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);
```

Copyright © 2013, Juniper Networks, Inc.

```
      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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    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;
    }
    (source-filtering | no-source-filtering);
  }
  flexible-vlan-tagging;
  gigether-options {
    802.3ad aex;
    (asynchronous-notification | no-asynchronous-notification);
    (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
      local-interface-offline>;
    auto-reconnect seconds;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    no-auto-mdix;
    source-address-filter {
      mac-address;
    }
    (source-filtering | no-source-filtering);
    ethernet-switch-profile {
      (mac-learn-enable | no-mac-learn-enable);
      tag-protocol-id [ tpids ];
      ethernet-policer-profile {
```

```
        input-priority-map {
          ieee802.1p premium [ values ];
        }
        output-priority-map {
          classifier {
            premium {
              forwarding-class class-name {
                loss-priority (high | low);
              }
            }
          }
        }
        policer cos-policer-name {
          aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
          }
          premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
          }
        }
      }
    }
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  ima-group-options {
    differential-delay number;
    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);
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
}
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 | ni1 | 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 {
```

```
        interfaceinterface-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 value;
      (no-start-measurement | 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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
}
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-vcipsn-vci-identifier;
psn-vpipsn-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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
      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);
```

Copyright © 2013, Juniper Networks, Inc.

```
    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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
  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;
```


Copyright © 2013, Juniper Networks, Inc.

```
        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;
        }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
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-vsa-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>;
```

 Copyright © 2013, Juniper Networks, Inc.

```
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 ];
  }
}
```

Copyright © 2013, Juniper Networks, Inc.

```
            }
          }
        }
      }
```

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;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    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;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
}
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;
```

Copyright © 2013, Juniper Networks, Inc.

```
              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*

[edit protocols connections] Hierarchy Level

The following statements can also be configured at the **[edit logical-systems *logical-system-name* protocols connections]** hierarchy level.

```
interface-switch connection-name {
```

 Copyright © 2013, Juniper Networks, Inc.

```
    interface interface-name.unit-number;
    interface interface-name.unit-number;
}
```

Related Documentation

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

[edit protocols dot1x] Hierarchy Level

```
dot1x {
  authenticator
    authentication-profile-name access-profile-name;
    interface interface-ids {
      maximum-requests integer;
      retries integer;
      quiet-period seconds;
      transmit-period seconds;
      reauthentication (disable | interval seconds);
      server-timeout seconds;
      supplicant (single);
      supplicant-timeout seconds;
    }
  }
}
```

Related Documentation

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

[edit protocols iccp] Hierarchy Level

```
iccp {
  traceoptions;
  local-ip-address ip address;
  session-establishment-hold-time value;
  authentication-key string;
  peer ip-address {
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    redundancy-group-id-list redundancy-group-id-list;
    liveness-detection;
  }
}
```

Related Documentation

- *traceoptions*
- *Configuring ICCP for MC-LAG*

[edit protocols lacp] Hierarchy Level

```
traceoptions {
```


Copyright © 2013, Juniper Networks, Inc.

```
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>);
  }
```

Related Documentation

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

[edit protocols lldp] Hierarchy Level

```
protocols {
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier seconds;
    interface (all | interface-name) {
      disable;
    }
    lldp-configuration-notification-interval seconds;
    ptopo-configuration-maximum-hold-time seconds;
    ptopo-configuration-trap-interval seconds;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    transmit-delay seconds;
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols oam] Hierarchy Level

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
      event {
        adjacency-loss;
        interface-status-tlv (down | lower-layer-down);
        port-status-tlv blocked;
        rdi;
      }
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
```

Copyright © 2013, Juniper Networks, Inc.

```
}
maintenance-domain domain-name {
  bridge-domain name;
  routing-instance r1 {
  bridge-domain name;
  instance vpls-instance;
  interface (ge | xe) fpc/pic/port.domain;
  level number;
  maintenance-association name{
    mep identifier {
      direction (up | down)
      interface (ge | xe) fpc/pic/port.domain (working | protect );
      auto-discovery;
      lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
        rem-err-xcon | xcon);
      priority number;
    }
  }
  mip-half-function (none | default | explicit);
  name-format (character-string | none | dns | mac+2oct);
  short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    protect-maintenance-association protect-ma-name;
    remote-maintenance-association remote-ma-name;
    continuity-check {
      hold-interval minutes;
      interval (10m | 10s | 1m | 1s | 100ms);
      loss-threshold number;
    }
    maintenance-association ma-name {
    mip-half-function (none | default | explicit);
    mep mep-id {
      auto-discovery;
      direction (up | down);
      interface interface-name (working |protect);
      priority number;
      remote-mep mep-id {
        action-profile profile-name;
        sla-iterator-profile profile-name {
          data-tlv-size bytes;
          iteration-count frames;
          priority priority-value;
        }
      }
    }
  }
}
performance-monitoring {
  hardware-assisted-timestamping;
  sla-iterator-profiles {
    profile-name {
      disable;
      calculation-weight {
        delay delay-weight;
        delay-variation delay-variation-weight;
      }
      cycle-time milliseconds;
```


Copyright © 2013, Juniper Networks, Inc.

```
                iteration-period connections;
                measurement-type (loss | statistical-frame-loss | two-way-delay);
              }
            }
          }
        }
        link-fault-management {
          action-profile profile-name {
            action {
              syslog;
              link-down;
              send-critical-event;
            }
            event {
              link-adjacency-loss;
              link-event-rate {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
              protocol-down;
            }
          }
          interface interface-name {
            apply-action-profile profile-name;
            event-thresholds {
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            link-discovery (active | passive);
            negotiation-options {
              allow-remote-loopback;
              no-allow-link-events;
            }
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
          }
        }
        fnp {
          interval <100ms | 1s | 10s | 1m | 10m>;
          loss-threshold number
          interface interface name {
              domain-id domain-id
          }
        }
      }
```

Related Documentation

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

Copyright © 2013, Juniper Networks, Inc.

[edit protocols ppp] Hierarchy Level

```
monitor-session (interface-name | all);
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable> ;
  flag flag <disable>;
}
```

Related Documentation

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

[edit protocols pppoe] Hierarchy Level

```
protocols {
  pppoe {
    no-send-pads-error;
    no-send-pads-ac-info
    pado-advertise;
    service-name-tables table-name {
      service service-name {
        agent-specifier {
          aci circuit-id-string ari remote-id-string {
            (delay seconds | drop | terminate);
            dynamic-profile profile-name;
            routing-instance routing-instance-name;
            static-interface interface-name;
          }
        }
        (delay seconds | drop | terminate);
        dynamic-profile profile-name;
        max-sessions number;
        routing-instance routing-instance-name;
      }
    }
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      filter {
        aci regular-expression;
        ari regular-expresion;
        service-name regular-expresion;
        underlying-interface interface-name;
      }
      flag flag;
      level (all | error | info | notice | verbose | warning);
      no-remote-trace;
    }
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

[edit protocols protection-group] Hierarchy Level

```
ethernet-ringring-name {
  east-interface {
    control-channel channel-name {
      vlan number;
    }
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
  west-interface {
    control-channel channel-name {
      vlan number;
    }
  }
}
```

Related Documentation

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

[edit protocols vrrp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  vrrp {
    asymmetric-hold-time;
    delegate-processing;
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
```

```
            version-3;
        }
    }
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

PART 2

Configuring Ethernet Interfaces

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 2

Configuring Ethernet Interfaces

You can configure the following properties specific to aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces:

Ethernet Interfaces Overview

Ethernet was developed in the early 1970s at the Xerox Palo Alto Research Center (PARC) as a data-link control layer protocol for interconnecting computers. It was first widely used at 10 megabits per second (Mbps) over coaxial cables and later over unshielded twisted pairs using 10Base-T. More recently, 100Base-TX (Fast Ethernet, 100 Mbps), Gigabit Ethernet (1 gigabit per second [Gbps]), 10-Gigabit Ethernet (10 Gbps), and 100-Gigabit Ethernet (100 Gbps) have become available.

Juniper Networks routers support the following types of Ethernet interfaces:

- Fast Ethernet
- Tri-Rate Ethernet copper

- Gigabit Ethernet
- Gigabit Ethernet intelligent queuing (IQ)
- Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet
- 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM)
- 100-Gigabit Ethernet
- Management Ethernet interface, which is an out-of-band management interface within the router
- Internal Ethernet interface, which connects the Routing Engine to the packet forwarding components
- Aggregated Ethernet interface, a logical linkage of Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet physical connections

Related Documentation

- Configuring Ethernet Physical Interface Properties on page 36
- Configuring J Series Services Router Switching Interfaces on page 40
- MX Series Router Interface Identifiers on page 42
- Enabling Ethernet MAC Address Filtering on page 42
- Configuring Ethernet Loopback Capability on page 45
- Configuring Flow Control on page 45
- Ignoring Layer 3 Incomplete Errors on page 46
- Configuring the Link Characteristics on Ethernet Interfaces on page 47
- Configuring Gratuitous ARP on page 48
- Adjusting the ARP Aging Timer on page 49
- Configuring the Interface Speed on Ethernet Interfaces on page 49
- Configuring the Ingress Rate Limit on page 50
- Configuring Multicast Statistics Collection on Ethernet Interfaces on page 51
- Configuring Weighted Random Early Detection on page 51
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

Configuring Ethernet Physical Interface Properties

To configure Fast Ethernet-specific physical interface properties, include the **fastether-options** statement at the **[edit interfaces fe-*fpc/pic/port*]** hierarchy level:

[edit interfaces fe-*fpc/pic/port*]

 Copyright © 2013, Juniper Networks, Inc.

```
link-mode (full-duplex | half-duplex);
speed (10m | 100m);
vlan-tagging;
fastether-options {
  802.3ad aex (primary | backup);
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  ingress-rate-limit rate;
  (loopback | no-loopback);
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

NOTE: The **speed** statement applies to the management Ethernet interface (**fxp0** or **em0**), the Fast Ethernet 12-port and 48-port Physical Interface Card (PIC) interfaces, the J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces. The Fast Ethernet, **fxp0**, and **em0** interfaces can be configured for 10 Mbps or 100 Mbps (**10m | 100m**). The J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps (**10m | 100m | 1g**). The 4-port and 8-port Fast Ethernet PICs support a speed of 100 Mbps only.

MX Series routers support Gigabit Ethernet automatic line sensing of MDI (Media Dependent Interface) and MDIX (Media Dependent Interface with Crossover) port connections. MDI is the Ethernet port connection typically used on network interface cards (NIC). MDIX is the standard Ethernet port wiring for hubs and switches. This feature allows MX Series routers to automatically detect MDI and MDIX connections and configure the router port accordingly. You can disable this feature by using the **no-auto-mdix** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level.

NOTE: Junos OS supports Ethernet host addresses with no subnets. This enables you to configure an Ethernet interface as a host address (that is, with a network mask of /32), without requiring a subnet. Such interfaces can serve as OSPF point-to-point interfaces, and MPLS is also supported.

To configure physical interface properties specific to Gigabit Ethernet and 10-Gigabit Ethernet, include the **gigether-options** statement at the **[edit interfaces ge-*fpc/pic/port*]** or **[edit interfaces xe-*fpc/pic/port*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
gigether-options {
  802.3ad aex (primary | backup);
  auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  (flow-control | no-flow-control);
```

```
    ignore-l3-incompletes;
    (loopback | no-loopback);
    no-auto-mdix;
    source-address-filter {
      mac-address;
    }
    (source-filtering | no-source-filtering);
  }
```

Additionally, for 10-Gigabit Ethernet DWDM-specific physical interface properties, include the **optics-options** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
optics-options {
  wavelength nm;
}
```

To configure Gigabit Ethernet IQ-specific physical interface properties, include the **gigether-options** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level. These statements are supported on 10-Gigabit Ethernet IQ2 and IQ2-E PIC. Some of these statements are also supported on Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router). For more information, see "Example: Configuring Gigabit Ethernet Interfaces" on page 543.

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
gigether-options {
  802.3ad aex (primary | backup);
  auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  (source-filtering | no-source-filtering);
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [tpids ];
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
```

 Copyright © 2013, Juniper Networks, Inc.

```
                premium {
                    bandwidth-limit bps;
                    burst-size-limit bytes;
                }
            }
        }
    }
    native-vlan-id number;
}
```

To configure 10-Gigabit Ethernet physical interface properties, include the **lan-phy** or **wan-phy** statement at the **[edit interfaces xe-*fpc/pic/port* framing]** hierarchy level. For more information, see "10-Gigabit Ethernet Framing Overview" on page 443.

```
[edit interfaces]
xe-0/0/0 {
    framing {
        (lan-phy | wan-phy);
    }
}
```

To configure OAM 802.3ah support for Ethernet interfaces, include the **oam** statement at the **[edit protocols]** hierarchy level.

```
oam {
    ethernet {
        link-fault-management {
            interfaces {
                interface-name{
                    pdu-interval interval;
                    link-discovery (active | passive);
                    pdu-threshold count;
                }
            }
        }
    }
}
```

To configure Gigabit Ethernet IQ-specific logical interface properties, include the **input-vlan-map**, **output-vlan-map**, **layer2-policer**, and **vlan-tags** statements:

```
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;
}
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;
}
layer2-policer {
```

```
        input-policer policer-name;
        input-three-color policer-name;
        output-policer policer-name;
        output-three-color policer-name;
    }
    vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

To configure aggregated Ethernet-specific physical interface properties, include the **aggregated-ether-options** statement at the **[edit interfaces ae*x*]** hierarchy level:

```
[edit interfaces aex]
aggregated-ether-options {
    ethernet-switch-profile {
        tag-protocol-id tpid;
    }
    (flow-control | no-flow-control);
    lacp mode {
        periodic interval;
    }
    link-protection;
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
```

Related Documentation

- Example: Configuring Gigabit Ethernet Interfaces on page 543
- 10-Gigabit Ethernet Framing Overview on page 443
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Configuring J Series Services Router Switching Interfaces

The J Series routers with multiport Gigabit Ethernet uPIMs supports Ethernet access switching. This functionality provides the ability to switch traffic at Layer 2 in addition to routing traffic at Layer 3.

J Series routers with multiport Gigabit Ethernet uPIMs can be deployed in branch offices as an access or desktop switch with integrated routing capability. The multiport Gigabit Ethernet uPIM provides Ethernet switching, while the Routing Engine provides routing functionality.

 Copyright © 2013, Juniper Networks, Inc.

Routed traffic is forwarded from any port of the multiport Gigabit Ethernet uPIM to the WAN interface. Switched traffic is forwarded from one port of the multiport Gigabit Ethernet uPIM to another port on the same the multiport Gigabit Ethernet uPIM. Switched traffic is not forwarded from a port on one multiport Gigabit Ethernet uPIM to a port on a different multiport Gigabit Ethernet uPIM. For more information about configuring the multiport Gigabit Ethernet uPIM switching mode, see the *Junos OS Administration Library for Routing Devices*.

In access switching mode, only one physical interface is configured for the entire multiport Gigabit Ethernet uPIM. The single physical interface serves as a Virtual Router Interface (VRI). Configuration of the physical port characteristics is done under the single physical interface.

To configure multiport Gigabit Ethernet uPIM Ethernet port properties, include the **switch-port** statement at the **[edit interfaces ge-*pim*/0/0]** hierarchy level:

```
[edit interfaces ge-pim/0/0]
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed 1g;
    link-mode (full-duplex | half-duplex);
  }
}
```

Access switching mode is supported on the 6-port, 8-port, and 16-port Gigabit Ethernet uPIMs.

The multiport Gigabit Ethernet uPIMs are supported on the J2320, J2350, J4350, and J6350 Services Routers.

The 6-port and 8-port multiport Gigabit Ethernet uPIM occupies a single slot and can be installed in any slot. Because the 16-port Gigabit Ethernet uPIM is two slots high, you cannot install a 16-port uPIM in the top slots (slots 1 and 4). Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

Example: Configuring J Series Services Router Switching Interfaces

Configure a single physical interface for the uPIM and set the port parameters for port 0 and port 1:

```
[edit interfaces]
ge-2/0/0 {
  switch-options {
    switch-port 0 {
      no-auto-negotiation;
      speed 1g;
      link-mode full-duplex;
    }
    switch-port 1 {
      no-auto-negotiation;
      speed 10m;
      link-mode half-duplex;
```

Copyright © 2013, Juniper Networks, Inc.

```
      }
    }
  }
```

Related Documentation

- switch-options on page 656
- switch-port on page 657
- speed on page 650
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

MX Series Router Interface Identifiers

Juniper Networks MX Series 3D Universal Edge Routers support several types of line cards, including Dense Port Concentrators (DPCs), Flexible Port Concentrators (FPCs) with associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) with associated Modular Interface Cards (MICs), or MICs. FPCs are populated with PICs for various interface types. DPCs and MPCs with associated MICs, and MICs support a variety of port configurations and combine the functions of FPCs and the PICs. The configuration syntax for each type of line card is the same: *type-fpc/pic/port*.

Ports are numbered from 0 through 9 for Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. Port numbers are always 0 for 10-Gigabit Ethernet interfaces.

NOTE: In certain displays, the MX Series routers identify the Packet Forwarding Engine (PFE) rather than the PIC number. PFE 0 corresponds to PIC 0, PFE 1 corresponds to PIC 2, PFE 2 corresponds to PIC 1, and PFE 3 corresponds to PIC 3.

Related Documentation

- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*
- *Router Interfaces*

Enabling Ethernet MAC Address Filtering

By default, source address filtering is disabled. On aggregated Ethernet interfaces, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can enable source address filtering, which blocks all incoming packets to an interface.

NOTE: Source address filtering is not supported on J Series Services Routers.

 Copyright © 2013, Juniper Networks, Inc.

To enable the filtering, include the **source-filtering** statement:

```
source-filtering;
```

To explicitly disable filtering, include the **no-source-filtering** statement:

```
no-source-filtering;
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

NOTE: When you integrate a standalone T640 router into a routing matrix, the PIC media access control (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, when you integrate a T1600 or T4000 router into a routing matrix, the PIC MAC addresses for the integrated T1600 or T4000 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 or T4000 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

Filtering Specific MAC Addresses

When source address filtering is enabled, you can configure the interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the **source-address-filter** statement:

```
source-address-filter {
  mac-address;
  <additional-mac-address>;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

You can specify the MAC address as ***nn*:*nn*:*nn*:*nn*:*nn*:*nn*** or ***nnnn* .*nnnn*.*nnnn***, where **n** is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include the **source-address-filter** statement multiple times.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The **source-address-filter** statement is not supported on Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router); instead, include the **accept-source-mac** statement. For more information, see "Configuring MAC Address Filtering" on page 401.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

Source address filtering does not work when Link Aggregation Control Protocol (LACP) is enabled. This behavior is not applicable to T series routers and PTX Series Packet Transport Routers. For more information about LACP, see "Configuring Aggregated Ethernet LACP" on page 177.

NOTE: On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement at the [edit interfaces **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

NOTE: The **source-address-filter** statement is not supported on MX Series routers with MPC4E (model numbers: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE); instead, include the **accept-source-mac** statement. For more information, see "Configuring MAC Address Filtering" on page 401.

Related Documentation

- **source-address-filter on page 647**
- Configuring MAC Address Filtering on page 401
- Configuring Aggregated Ethernet LACP on page 177
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Configuring Ethernet Loopback Capability

By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

NOTE: If you configure a local loopback on a 1-port 10-Gigabit IQ2 and IQ2-E PIC using the **loopback** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level, the transmit-path stops working, causing the remote end to detect a link down.

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

Related Documentation

- loopback on page 600
- Ethernet Interfaces Overview on page 35
- *EX Series Switches Interfaces Overview*
- *Ethernet Interfaces*

Configuring Flow Control

By default, the router or switch imposes flow control to regulate the amount of traffic sent out on a Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interface. Flow control is not supported on the 4-port Fast Ethernet PIC. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the router or switch to permit unrestricted traffic. To disable flow control, include the **no-flow-control** statement:

```
no-flow-control;
```

To explicitly reinstate flow control, include the **flow-control** statement:

```
flow-control;
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Related Documentation

- flow-control on page 573
- Ethernet Interfaces Overview on page 35
- *EX Series Switches Interfaces Overview*
- *Ethernet Interfaces*

Ignoring Layer 3 Incomplete Errors

By default, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces count Layer 3 incomplete errors. You can configure the interface to ignore Layer 3 incomplete errors.

To ignore Layer 3 incomplete errors, include the **ignore-l3-incompletes** statement:

```
ignore-l3-incompletes;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

Related Documentation

- ignore-l3-incompletes on page 580
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Configuring the Link Characteristics on Ethernet Interfaces

Full-duplex communication means that both ends of the communication can send and receive signals at the same time. *Half-duplex* is also bidirectional communication, but signals can flow in only one direction at a time.

By default, the router's management Ethernet interface, **fxp0** or **em0**, autonegotiates whether to operate in full-duplex or half-duplex mode. J Series Gigabit Ethernet interfaces and Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet and 10-Gigabit Ethernet, the link partner must also be set to full duplex.

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

NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: On a J Series ePIM Fast Ethernet interface, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."

NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

NOTE: Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.

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

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

Related Documentation

- link-mode on page 593
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests provide duplicate IP address detection. A gratuitous ARP request is a broadcast request for a router's own IP address. If a router or switch sends an ARP request for its own IP address and no ARP replies are received, the router- or switch-assigned IP address is not being used by other nodes. If a router or switch sends an ARP request for its own IP address and an ARP reply is received, the router- or switch-assigned IP address is already being used by another node.

By default, the router or switch responds to gratuitous ARP requests. On Ethernet interfaces, you can disable responses to gratuitous ARP requests. To disable responses to gratuitous ARP requests, include the **no-gratuitous-arp-request** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
no-gratuitous-arp-request;
```

To return to the default—that is, to respond to gratuitous ARP requests—delete the **no-gratuitous-arp-request** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name no-gratuitous-arp-request
```

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router or switch receives a gratuitous ARP reply, the router or switch can insert an entry for that reply in the ARP cache.

By default, updating the ARP cache on gratuitous ARP replies is disabled on the router or switch. On Ethernet interfaces, you can enable handling of gratuitous ARP replies on a specific interface by including the **gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

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

 Copyright © 2013, Juniper Networks, Inc.

gratuitous-arp-reply;

To restore the default behavior, include the **no-gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

[edit interfaces *interface-name*]
no-gratuitous-arp-reply;

Related Documentation

- gratuitous-arp-reply on page 578
- no-gratuitous-arp-request on page 609
- Ethernet Interfaces Overview on page 35
- *EX Series Switches Interfaces Overview*
- *Ethernet Interfaces*

Adjusting the ARP Aging Timer

By default, the ARP aging timer is set at 20 minutes. In environments with many directly attached hosts, such as metro Ethernet environments, increasing the amount of time between ARP updates by configuring the ARP aging timer can improve performance in an event where having thousands of clients time out at the same time might impact packet forwarding performance. In environments where there are devices connected with lower ARP aging timers (less than 20 minutes), decreasing the ARP aging timer can improve performance by preventing the flooding of traffic toward next hops with expired ARP entries. In most environments, the default ARP aging timer value does not need to be adjusted.

To configure the system-wide ARP aging timer, include the **aging-timer** statement at the **[edit system arp]** hierarchy level:

[edit system arp]
user@host# **aging-timer *minutes***

The aging timer range is from 1 through 240 minutes. The timer value you configure takes effect as ARP entries expire. In other words, each subsequent refreshed ARP entry receives the new timer value. The new timer value does not apply to ARP entries that exist at the time you commit the configuration.

For more information about statements you can configure at the **[edit system]** hierarchy level, see the *Junos OS Administration Library for Routing Devices*.

Related Documentation

- *arp (System)*
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Configuring the Interface Speed on Ethernet Interfaces

For M Series and T Series Fast Ethernet 12-port and 48-port PIC interfaces, the management Ethernet interface (**fxp0** or **em0**), the J Series Gigabit Ethernet uPIM interfaces, and the MX Series Tri-Rate Ethernet copper interfaces, you can explicitly set the interface speed. The Fast Ethernet, **fxp0**, and **em0** interfaces can be configured for 10 Mbps or 100 Mbps **(10m | 100m)**. The J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps **(10m | 100m | 1g)**. MX Series routers, with MX-DPC and Tri-Rate Copper SFPs, support 20x1 Copper to provide backwards compatibility with 100/10BASE-T and 1000BASE-T operation through an Serial Gigabit Media Independent Interface (SGMII) interface.

NOTE: On MX Series routers with tri-rate copper SFP interfaces, if the port speed is negotiated to the configured value and the negotiated speed and interface speed do not match, the link will not be brought up.

NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: Half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the **speed** statement, you must include the **link-mode full-duplex** statement at the same hierarchy level.

To explicitly configure the speed, include the **speed** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
speed (10m | 100m | 1g);
```

Related Documentation

- speed on page 650
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Configuring the Ingress Rate Limit

On Fast Ethernet 8-port, 12-port, and 48-port PIC interfaces only, you can apply port-based rate limiting to the ingress traffic that arrives at the PIC.

To configure an ingress rate limit on a Fast Ethernet 8-port, 12-port, or 48-port PIC interface, include the **ingress-rate-limit** statement at the **[edit interfaces *interface-name* fastether-options]** hierarchy level:

```
[edit interfaces interface-name fastether-options]
ingress-rate-limit rate;
```

rate can range in value from 1 through 100 Mbps.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- **ingress-rate-limit on page 580**
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Configuring Multicast Statistics Collection on Ethernet Interfaces

T Series and TX Matrix routers support multicast statistics collection on Ethernet interfaces in both ingress and egress directions. The multicast statistics functionality can be configured on a physical interface thus enabling multicast accounting for all the logical interfaces below the physical interface.

The multicast statistics information is displayed only when the interface is configured with the **multicast-statistics** statement, which is not enabled by default.

Multicast statistics collection requires at least one logical interface is configured with family inet and/or inet6; otherwise, the commit for **multicast-statistics** will fail.

The multicast in/out statistics can be obtained via interfaces statistics query through CLI and via MIB objects through SNMP query.

To configure multicast statistics:

1. Include the **multicast-statistics** statement at the **[edit interfaces interface-name]** hierarchy level.

An example of a multicast statistics configuration for a Ethernet interface follows:

```
[edit interfaces]
  ge-fpc/pic/port {
  multicast-statistics;
}
```

To display multicast statistics, use the **show interfaces *interface-name* statistics detail** command.

Related Documentation

- *multicast-statistics*
- Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces on page 188
- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Configuring Weighted Random Early Detection

On M7i, M10i, M40e, M320, M120, and T Series routers, the Ethernet IQ2 and IQ2-E PIC families extend CoS functionality by supporting network congestion avoidance with weighted random early detection (WRED).

Related Documentation

- For information on configuring WRED, see the *Junos OS Class of Service Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

- Ethernet Interfaces Overview on page 35
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 3

Configuring 802.1Q VLANs

802.1Q VLANs Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Related Documentation

- Configuring Dynamic 802.1Q VLANs on page 54
- 802.1Q VLAN IDs and Ethernet Interface Types on page 55

Configuring Dynamic 802.1Q VLANs

You can configure the router to dynamically create VLANs when a client accesses an interface and requests a VLAN ID that does not yet exist. When a client accesses a VLAN interface, the router instantiates a VLAN dynamic profile that you have associated with the interface. Using the settings in the dynamic profile, the router extracts information about the client from the incoming packet (for example, the interface and unit values), saves this information in the routing table, and creates a VLAN or stacked VLAN ID for the client from a range of VLAN IDs that you configure for the interface.

 Copyright © 2013, Juniper Networks, Inc.

Dynamically configuring VLANs or stacked VLANs requires the following general steps:

1. Configure a dynamic profile for dynamic VLAN or dynamic stacked VLAN creation.
2. Associate the VLAN or stacked VLAN dynamic profile with the interface.
3. Specify the Ethernet packet type that the VLAN dynamic profile accepts.
4. Define VLAN ranges for use by the dynamic profile when creating VLAN IDs.

For procedures on how to configure dynamic VLANs and dynamic stacked VLANs for client access, see the *Junos OS Subscriber Management and Services Library*.

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

802.1Q VLAN IDs and Ethernet Interface Types

You can partition the router into up to 4095 different VLANs—depending on the router model and the physical interface types—by associating logical interfaces with specific VLAN IDs.

VLAN ID 0 is reserved for tagging the priority of frames. VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN circuit cross-connect (CCCs).

For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. With flexible Ethernet services encapsulation, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

The maximum number of user-configurable VLANs is 15 on each port of the Dense-FE PIC (8-port/12-port/48-port).

Table 3 on page 55 lists VLAN ID range by interface type.

Table 3: VLAN ID Range by Interface Type

Interface Type	VLAN ID Range
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
4-port, 8-port, and 12-port Fast Ethernet	1 through 1023
48-port Fast Ethernet	1 through 4094
Tri-Rate Ethernet copper	1 through 4094
Gigabit Ethernet	1 through 4094

Copyright © 2013, Juniper Networks, Inc.

Table 3: VLAN ID Range by Interface Type *(continued)*

Interface Type	VLAN ID Range
Gigabit Ethernet IQ	1 through 4094
10-Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023

NOTE: For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the built-in Gigabit Ethernet port on the M7i router), VLAN IDs on a single interface can differ from each other.

Because IS-IS has an 8-bit limit for broadcast multiaccess media, you cannot set up more than 255 adjacencies over Gigabit Ethernet using VLAN tagging. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Enabling VLAN Tagging

You can configure the router to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames. For more information, see the following sections:

NOTE: If you configure VLAN tagging on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces on M320, M120, and T Series routers, the Junos OS creates an internal logical interface that reserves 50 Kbps of bandwidth from Gigabit Ethernet IQ interfaces and 2 Mbps of bandwidth from Gigabit Ethernet IQ2 and IQ2-E interfaces. As a result, the effective available bandwidth for these interface types is now 999.5 Mbps and 998 Mbps, respectively.

Copyright © 2013, Juniper Networks, Inc.

Configuring Single-Tag Framing

To configure the router to receive and forward single-tag frames with 802.1Q VLAN tags, include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
vlan-tagging;
```

Configuring Dual Tagging

To configure the routing platform to receive and forward dual-tag frames with 802.1Q VLAN tags, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
stacked-vlan-tagging;
```

Configuring Mixed Tagging

Mixed tagging is supported for Gigabit Ethernet interfaces on Gigabit Ethernet IQ2 and IQ2-E, and IQ or IQE PICs on M Series and T Series routers, for all MX Series router Gigabit and 10-Gigabit Ethernet interfaces, and for aggregated Ethernet interfaces with member links in IQ2 and IQ2-E PICs or in MX Series DPCs. Mixed tagging lets you configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

NOTE: Mixed tagging is not supported on Fast Ethernet interfaces or on J Series Services Routers.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level. You must also include the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
  vlan-id number;
  family family {
    address address;
  }
}
unit logical-unit-number {
  vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
  family family {
    address address;
  }
}
```

NOTE: When you configure the physical interface MTU for mixed tagging, you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

If the same physical interface MTU value is configured on both the VLAN and flexible VLAN-tag routers, the L2 circuit configuration does not come up and a MTU mismatch is logged. However, normal traffic flow is unaffected.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid. See "Configuring VLAN Encapsulation" on page 65.

Configuring Mixed Tagging Support for Untagged Packets

For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs, for all MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces configured for 802.1Q flexible VLAN tagging, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, you can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the **native-vlan-id** statement and the **flexible-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;
```

The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Example: Configuring Mixed Tagging

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

```
[edit interfaces ge-3/0/1]
flexible-vlan-tagging;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
```

 Copyright © 2013, Juniper Networks, Inc.

```
    family inet {
      address 10.66.1.2/30;
    }
  }
```

For information about binding VLAN IDs to logical interfaces, see “Binding VLAN IDs to Logical Interfaces” on page 59. For information about configuring dual VLAN tags using the **vlan-tag** statement, see “Stacking a VLAN Tag” on page 206.

Example: Configuring Mixed Tagging to Support Untagged Packets

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Binding VLAN IDs to Logical Interfaces

The following sections describe how to configure logical interfaces to receive and forward VLAN-tagged frames:

Binding VLAN IDs to Logical Interfaces Overview

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, a range of VLAN IDs, or a list of VLAN IDs to the logical interface. Table 4 on page 60 lists the configuration statements you use to bind VLAN IDs to logical interfaces, organized by scope of the VLAN IDs used to match incoming packets:

Table 4: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces

Scope of VLAN ID Matching	Type of VLAN Framing Supported on the Logical Interface	
	Single-Tag Framing	Dual-Tag Framing
VLAN ID	**vlan-id** *vlan-id*;	**vlan-tags outer** *tpid*.<*vlan-id*> **inner** *tpidvlan-id*;
VLAN ID Range	**vlan-id-range** *vlan-id–vlan-id*;	**vlan-tags outer** *tpid.vlan-id* **inner-range** *tpid.vlan-id–vlan-id*;
VLAN ID List	**vlan-id-list** [*vlan-id vlan-id–vlan-id*];	**vlan-tags outer** <*tpid*.>*vlan-id* **inner-list** [*vlan-id vlan-id–vlan-id*];

You can include all of the statements at the following hierarchy levels:

- **[edit interfaces** *interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number***]**

NOTE: The **inner-list** option of the **vlan-tags** statement does not support Tag Protocol ID (TPID) values.

Binding a VLAN ID to a Logical Interface

A logical interface that you have associated (bound) to a particular VLAN ID will receive and forward incoming frames that contain a matching VLAN ID.

Binding a VLAN ID to a Single-Tag Logical Interface

To bind a VLAN ID to a single-tag logical interface, include the **vlan-id** statement:

```
vlan-id vlan-id;
```

You can include the statement at the following hierarchy levels:

- **[edit interfaces** *ethernet-interface-name* **unit** *logical-unit-number***]**
- **[edit logical-systems** *logical-system-name* **interfaces** *ethernet-interface-name* **unit** *logical-unit-number***]**

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces** *ethernet-interface-name***]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a VLAN ID to a Dual-Tag Logical Interface

To bind a VLAN ID to a dual-tag logical interface, include the **vlan-tags** statement:

```
vlan-tags inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

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

 Copyright © 2013, Juniper Networks, Inc.

- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a Range of VLAN IDs to a Logical Interface

A VLAN range can be used by service providers to interconnect multiple VLANs belonging to a particular customer over multiple sites. Using a VLAN ID range conserves switch resources and simplifies configuration.

Binding a Range of VLAN IDs to a Single-Tag Logical Interface

To bind a range of VLAN IDs to a single-tag logical interface, include the **vlan-id-range** statement:

vlan-id-range *vlan-id-vlan-id*;

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a Range of VLAN IDs to a Dual-Tag Logical Interface

To bind a range of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner–list** option to specify the VLAN IDs as an inclusive range by separating the starting VLAN ID and ending VLAN ID with a hyphen.

vlan-tags inner–list [*vlan-id vlan-id–vlan-id*] outer <*tpid.*>*vlan-id*;

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Example: Binding Ranges VLAN IDs to Logical Interfaces

The following example configures two different ranges of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-3/0/0 {
```

```
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 500-600;
  }
}
ge-3/0/1 {
  flexible-vlan-tagging;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 200-300;
    }
    unit 1 {
      encapsulation vlan-bridge;
      vlan-tags outer 1000 inner-range 100-200;
    }
}
```

Binding a List of VLAN IDs to a Logical Interface

In Junos OS Release 9.5 and later, on MX Series routers and in Junos OS Release 12.2R2 and later on EX Series switches, you can bind a list of VLAN IDs to a single logical interface, eliminating the need to configure a separate logical interface for every VLAN or VLAN range. A logical interface that accepts packets tagged with any VLAN ID specified in a VLAN ID list is called a *VLAN-bundled* logical interface.

You can use VLAN-bundled logical interfaces to configure circuit cross-connects between Layer 2 VPN routing instances or Layer 2 circuits. Using VLAN-bundled logical interfaces simplifies configuration and reduces use of system resources such as logical interfaces, next hops, and circuits.

As an alternative to configuring multiple logical interfaces (one for each VLAN ID and one for each range of VLAN IDs), you can configure a single VLAN-bundled logical interface based on a list of VLAN IDs.

NOTE: The **vlan-id** option is not supported to achieve VLAN normalization on VPLS instances that are configured with **vlan-id-list**. However, you can use the **vlan-maps** option to achieve VLAN normalization.

Binding a List of VLAN IDs to a Single-Tag Logical Interface

To bind a list of VLAN IDs to a single-tag logical interface, include the **vlan-id-list** statement. Specify the VLAN IDs in the list individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

vlan-id-list [*vlan-id vlan-id–vlan-id*];

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

Copyright © 2013, Juniper Networks, Inc.

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a List of VLAN IDs to a Dual-Tag Logical Interface

To bind a list of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner-list** option to specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both:

vlan-tags inner-list [*vlan-id vlan-id–vlan-id*] outer <*tpid*>*vlan-id*;

NOTE: The inner-list option of the vlan-tags statement does not support Tag Protocol ID (TPID) values.

You can include the statement at the following hierarchy levels:

- **[edit interfaces *ethernet-interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]**

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Example: Binding Lists of VLAN IDs to Logical Interfaces

The following example configures two different lists of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-1/1/0 {
  vlan-tagging; # Only for single-tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id-list [20 30–40 45];
  }
}
ge-1/1/1 {
  flexible-vlan-tagging; # Only for mixed tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id-list [1 10 20 30–40];
  }
  unit 20 {
    encapsulation vlan-ccc;
    vlan-tags outer 200 inner-list [50–60 80 90–100];
  }
}
```

In the example configuration above, **ge-1/1/0** supports single-tag logical interfaces, and **ge-1/1/1** supports mixed tagging. The single-tag logical interfaces **ge-1/1/0.10** and **ge-1/1/1.20** each bundle lists of VLAN IDs. The dual-tag logical interface **ge-1/1/1.20** bundles lists of inner VLAN IDs.

TIP: You can group a range of identical interfaces into an interface range and then apply a common configuration to that interface range. For example, in the above example configuration, both interfaces ge-1/1/0 and ge-1/1/1 have the same physical encapsulation type of **flexible-ethernet-services**. Thus you can define an interface range with the interfaces ge-1/1/0 and ge-1/1/1 as its members and apply the encapsulation type flexible-ethernet-services to that defined interface range. For more information about interface ranges, see *Configuring Interface Ranges*.

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Configuring Interface Ranges*
- *Ethernet Interfaces*

Associating VLAN IDs to VLAN Demux Interfaces

The following sections describe how to configure VLAN demux interfaces to receive and forward VLAN-tagged frames:

Associating VLAN IDs to VLAN Demux Interfaces Overview

To configure a VLAN demux interface to receive and forward VLAN-tagged frames, you must associate a VLAN ID or dual tagged (stacked) VLAN ID to the interface. Table 5 on page 64 shows the configuration statements you use to associate VLAN IDs to VLAN demux interfaces, depending on the VLAN tag framing you use:

Table 5: Configuration Statements Used to Associate VLAN IDs to VLAN Demux Interfaces

	Single-Tag Framing	Dual-Tag Framing
Statement Format	vlan-id *vlan-id*;	vlan-tags outer *tpid*.<*vlan-id*> inner *tpidvlan-id*;

You can include all of the statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]
- [edit interfaces *demux0* unit *logical-unit-number*]

Copyright © 2013, Juniper Networks, Inc.

Associating a VLAN ID to a VLAN Demux Interface

A VLAN demux interface that you have associated to a particular VLAN ID receives and forwards incoming frames that contain a matching VLAN ID. You can associate a VLAN ID to a single-tag logical interface or to a dual-tagged (stacked) logical interface.

Associating a VLAN ID to a Single-Tag VLAN Demux Interface

To associate a VLAN ID to a single-tag VLAN demux interface, include the **vlan-id** statement at the **[edit interfaces *demux0* unit *logical-unit-number*]** hierarchy level:

```
vlan-id vlan-id;
```

To configure an interface to support single-tag logical interfaces, you must also include the **vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Associating a VLAN ID to a Dual-Tag VLAN Demux Interface

To associate a VLAN ID to a dual-tag VLAN demux interface, include the **vlan-tags** statement at the **[edit interfaces *demux0* unit *logical-unit-number*]** hierarchy level:

```
vlan-tags inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

To configure an interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Configuring VLAN Encapsulation

Gigabit Ethernet IQ, Gigabit Ethernet PICs with small form-factor pluggable optics (SFPs), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use flexible Ethernet services, VLAN CCC, or VLAN virtual private LAN service (VPLS) encapsulation.

Aggregated Ethernet interfaces configured for VPLS can use Ethernet VPLS or VLAN VPLS.

To configure the encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **flexible-ethernet-services**, **vlan-ccc**, or **vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (flexible-ethernet-services | vlan-ccc | vlan-vpls);
```

To configure the encapsulation on an aggregated Ethernet interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **flexible-ethernet-services**, **ethernet-vpls**, or **vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (flexible-ethernet-services | ethernet-vpls | vlan-vpls);
```

Copyright © 2013, Juniper Networks, Inc.

Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.

In general, you configure an interface's encapsulation at the **[edit interfaces *interface-name*]** hierarchy level. However, for some encapsulation types, including flexible Ethernet services, Ethernet VLAN CCC and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement:

```
encapsulation (vlan-ccc | vlan-tcc | vlan-vpls);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 600;
  }
}
```

Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

```
interfaces ae0 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 0 {
    vlan-id 100;
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Configuring VPLS Interface Encapsulation*
- *Ethernet Interfaces*

Configuring Extended VLAN Encapsulation

Gigabit Ethernet, 4-port Fast Ethernet, MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, 10-Gigabit Ethernet, and aggregated Ethernet interfaces with VLAN tagging enabled can use extended VLAN CCC or VLAN VPLS, which allow 802.1Q tagging. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, specifying **extended-vlan-ccc** or **extended-vlan-vpls**:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | extended-vlan-vpls);
```

For extended VLAN CCC and extended VLAN VPLS encapsulation, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: For extended VLAN CCC, the VLAN IDs on ingress and egress interfaces must be the same. For back-to-back connections, all VLAN IDs must be the same.

Example: Configuring Extended VLAN Encapsulation on a Gigabit Ethernet Interface

Configure extended VLAN CCC encapsulation on Gigabit Ethernet ingress and egress interfaces:

```
interfaces ge-0/0/0 {
   vlan-tagging;
   encapsulation extended-vlan-ccc;
   unit 0 {
      vlan-id 2;
      family ccc;
   }
}
interfaces ge-1/0/0 {
   vlan-tagging;
   encapsulation extended-vlan-ccc;
   unit 0 {
      vlan-id 2;
      family ccc;
   }
}
```

Example: Configuring Extended VLAN Encapsulation on an Aggregated Ethernet Interface

Configure extended VLAN VPLS encapsulation on an aggregated Ethernet interface:

```
interfaces ae0 {
```

Copyright © 2013, Juniper Networks, Inc.

```
    vlan-tagging;
    encapsulation extended-vlan-vpls;
    unit 0 {
      vlan-id 100;
    }
  }
```

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs

For MX Series routers, you can bind a list of VLAN IDs to a logical interface, configure a Layer 2 VPN routing instance or Layer 2 circuit on the logical interface, and then use the logical interface to configure a circuit cross-connect (CCC) to another Layer 2 VPN routing instance or Layer 2 circuit.

A CCC allows you to configure transparent connections between two circuits so that packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. You configure a CCC by connecting circuit interfaces of the same type. For more information, see *Circuit and Translational Cross-Connects Overview*.

NOTE: **The Junos OS supports binding of Ethernet logical interfaces to lists of VLAN IDs on MX Series routers only. For all other routers, you can bind an Ethernet logical interface to only a single VLAN ID or to a single range of VLAN IDs.**

The following configuration guidelines apply to bundling lists of VLAN IDs to Ethernet logical interfaces used to configure CCCs:

- Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs on page 68
- Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs on page 69

Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs

To enable a physical interface to support VLAN-bundled logical interfaces that you will use to configure a CCC, you must specify one of the following physical link-layer encapsulation types as the value of the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | flexible-ethernet-services);
```

- **extended-vlan-ccc**—For Ethernet interfaces with standard TPID tagging.
- **flexible-ethernet-services**—For supported Gigabit Ethernet interfaces for which you want to configure multiple per-unit Ethernet encapsulations.

 Copyright © 2013, Juniper Networks, Inc.

For more information about configuring the encapsulation on a physical interface, see *Configuring Interface Encapsulation on Physical Interfaces*.

Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs

For VLAN-bundled logical interfaces that you use to configure a CCC, specific logical link-layer encapsulation types are used inside the circuits themselves.

Table 6 on page 69 describes the logical link-layer encapsulation types used within circuits connected using VLAN-bundled logical interfaces of the same type.

Table 6: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces

Encapsulation Inside the Circuit	Layer 2 Circuit Joined by Configuring an Interface-to-Interface CCC Connection	
	Layer 2 VPN Routing Instance	Layer 2 Circuit
Syntax	encapsulation-type (ethernet | ethernet-vlan);	encapsulation vlan-ccc;
Hierarchy Level	[edit routing-instances *routing-instance-name* protocols l2vpn], [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]	[edit interfaces *ethernet-interface-name* unit *logical-unit-number*], [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]
Usage Guidelines	See the *Junos OS VPNs Library for Routing Devices*.	See *Configuring Interface Encapsulation on Logical Interfaces*, *Circuit and Translational Cross-Connects Overview*, and *Defining the Encapsulation for Switching Cross-Connects*.
For a Single-Tag Logical Interface	The MX Series router automatically uses **ethernet** as the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN. **NOTE:** With **ethernet** encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.	Configure the MX Series router to use **vlan-ccc** as the logical link-layer encapsulation type.
For a Dual-Tag Logical Interface	Configure the MX Series router to use **ethernet-vlan** as the Layer 2 protocol to encapsulate incoming traffic. With **ethernet-vlan** encapsulation, circuit signal processing checks that the VLAN ID list is the same at both ends of the CCC connection. If a VLAN ID list mismatch is detected, you can view the error condition in the **show interfaces** command output.	The MX Series router automatically uses **vlan-ccc** as the logical link-layer encapsulation type, regardless of the value configured.

Related Documentation

- 802.1Q VLANs Overview on page 53
- Binding VLAN IDs to Logical Interfaces on page 59

- *Defining the Encapsulation for Switching Cross-Connects*

Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 VPN routing instance on a logical interface bound to a list of VLAN IDs.

Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement on a provider edge (PE) router:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      vlan-id-list [vlan-id vlan-id–vlan-id]; # For single-tag
      vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id–vlan-id]; # For dual-tag
    }
    . . .
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

Specifying the Interface Over Which VPN Traffic Travels to the CE Router

To configure a Layer 2 VPN routing instance on a PE router, include the **instance-type** statement and specify the value **l2vpn**. To specify an interface connected to the router, include the **interface** statement and specify the VLAN-bundled logical interface:

```
instance-type l2vpn;
interface logical-interface-name;
```

You can include the statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

 Copyright © 2013, Juniper Networks, Inc.

Specifying the Interface to Handle Traffic for a CCC

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the following statements:

```
protocols {
  l2vpn {
    (control-word | no-control-word);
    encapsulation-type (ethernet | ethernet-vlan);
    site site-name {
      site-identifier identifier;
      interface logical-interface-name { # VLAN-bundled logical interface
        ... interface-options ...
      }
    }
  }
}
```

You can include the statements at the same hierarchy level at which you include the **instance-type l2vpn** and **interface *logical-interface-name*** statements:

- **[edit routing-instances *routing-instance-name*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

To enable a Layer 2 VPN routing instance on a PE router, include the **l2vpn** statement. For more information, see the *Junos OS VPNs Library for Routing Devices*.

The **encapsulation-type** statement specifies the Layer 2 protocol used for traffic from the customer edge (CE) router. If the Layer 2 VPN routing instance is being connected to a single-tag Layer 2 circuit, specify **ethernet** as the encapsulation type. If the Layer 2 VPN routing instance is being connected to a dual-tag Layer 2 circuit, specify **ethernet-vlan** as the encapsulation type.

To specify the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the **interface** statement and specify the VLAN-bundled logical interface.

Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 circuit on a logical interface bound to a list of VLAN IDs.

- Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance on page 71
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 72

Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the **vlan-id-list** statement or the **vlan-tags** statement:

```
interfaces {
  ethernet-interface-name {
```

Copyright © 2013, Juniper Networks, Inc.

```
        vlan-tagging; # Support single- or dual-tag logical interfaces
        flexible-vlan-tagging; # Support mixed tagging
        encapsulation (extended-vlan-ccc | flexible-ethernet-services);
        unit logical-unit-number {
          encapsulation vlan–ccc; # Required for single-tag
          vlan-id-list [vlan-id vlan-id–vlan-id]; # For single-tag
          vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id–vlan-id]; # For dual-tag
        }
        . . .
      }
    }
```

You can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

For a single-tag logical interface, include the **encapsulation** statement and specify **vlan-ccc** so that CCC circuit encapsulation is used inside the Layer 2 circuit.

NOTE: In the case of a dual-tag logical interface, the Junos OS automatically uses the **vlan-ccc** encapsulation type.

Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the following statements:

```
l2circuit {
  neighbor address {
    interface logical-interface-name {
      virtual-circuit-id number;
      no-control-word;
    }
  }
}
```

You can include the statements at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

To enable a Layer 2 circuit, include the **l2circuit** statement.

To configure the router as a neighbor for a Layer 2 circuit, specify the neighbor address using the **neighbor** statement.

To specify the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the **interface** statement and specify the VLAN-bundled logical interface.

Copyright © 2013, Juniper Networks, Inc.

Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface **ge-1/0/5.0** bundles a list of VLAN IDs, and the logical interface **ge-1/1/1.0** supports IPv4 traffic using IP address 10.30.1.130 and can participate in an MPLS path.

```
[edit interfaces]
ge-1/0/5 {
   vlan-tagging;
   encapsulation extended-vlan-ccc;
   unit 0 { # VLAN-bundled logical interface
      vlan-id-list [513 516 520-525];
   }
}
   ge-1/1/1 {
      unit 0 {
         family inet {
            address 10.30.1.1/30;
         }
         family mpls;
      }
   }
```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance:

```
[edit protocols]
rsvp {
   interface all;
   interface lo0.0;
}
mpls {
   label-switched-path lsp {
      to 10.255.69.128;
   }
   interface all;
}
bgp {
   group g1 {
      type internal;
      local-address 10.255.69.96;
      family l2vpn {
         signaling;
      }
      neighbor 10.255.69.128;
   }
}
ospf {
   traffic-engineering;
   area 0.0.0.0 {
      interface lo0.0;
      interface ge-1/1/1.0;
   }
}
```

Copyright © 2013, Juniper Networks, Inc.

The following configuration shows that the VLAN-bundled logical interface is the interface over which VPN traffic travels to the CE router and handles traffic for a CCC to which the VPN connects.

```
[edit routing-instances]
red {
   instance-type l2vpn;
   interface ge-1/0/5.0; # VLAN-bundled logical interface
   route-distinguisher 10.255.69.96:100;
   vrf-target target:1:1;
   protocols {
      l2vpn {
         encapsulation-type ethernet; # For single-tag VLAN logical interface
         site CE_ultima {
            site-identifier 1;
            interface ge-1/0/5.0;
         }
      }
   }
}
```

NOTE: Because the VLAN-bundled logical interface supports single-tag frames, Ethernet is the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN.

However, with Ethernet encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface **ge-1/0/5.0** bundles a list of VLAN IDs, and the logical interface **ge-1/1/1.0** supports IPv4 traffic using IP address 10.30.1.1/30 and can participate in an MPLS path.

```
[edit interfaces]
ge-1/0/5 {
   vlan-tagging;
   encapsulation extended-vlan-ccc;
   unit 0 { # VLAN-bundled logical interface
      vlan-id-list [513 516 520-525];
   }
}
ge-1/1/1 {
   unit 0 {
      family inet {
         address 10.30.1.1/30;
      }
```

 Copyright © 2013, Juniper Networks, Inc.

```
        family mpls;
    }
}
```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance, and shows that the VLAN-bundled logical interface handles traffic for a CCC to which the Layer 2 circuit connects:

```
[edit protocols]
rsvp {
    interface all;
    interface lo0.0;
}
mpls {
    label-switched-path lsp {
        to 10.255.69.128;
    }
    interface all;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-1/1/1.0;
    }
}
ldp {
    interface ge-1/1/1.0;
    interface ge-1/0/5.0; # VLAN-bundled logical interface
    interface lo0.0;
}
l2circuit {
    neighbor 10.255.69.128 {
        interface ge-1/0/5.0 { # VLAN-bundled logical interface
            virtual-circuit-id 3;
            no-control-word;
        }
    }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Configuring a Logical Interface for Access Mode

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified bridge domain. A logical interface configured to accept untagged packets is called an *access interface* or *access port*. Access interface configuration is supported on MX Series routers only.

```
interface-mode access;
```

You can include this statement at the following hierarchy levels:

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

When an untagged or tagged packet is received on an access interface, the packet is accepted, the VLAN ID is added to the packet, and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

Example: Configuring a Logical Interface for Access Mode

The following example configures a logical interface as an access port with a VLAN ID of 20:

```
[edit interfaces ge-1/2/0]
unit 1 {
   family bridge {
      interface-mode access;
      vlan-id 20;
   }
}
```

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Configuring a Logical Interface for Trunk Mode

As an alternative to configuring a logical interface for each VLAN, enterprise network administrators can configure a single logical interface to accept untagged packets or packets tagged with any VLAN ID specified in a list of VLAN IDs. Using a VLAN ID list conserves switch resources and simplifies configuration. A logical interface configured to accept packets tagged with any VLAN ID specified in a list is called a *trunk interface* or *trunk port*. Trunk interface configuration is supported on MX Series routers only. Trunk interfaces support integrated routing and bridging (IRB).

To configure a logical interface to accept any packet tagged with a VLAN ID that matches the list of VLAN IDs, include the **interface-mode** statement and specify the **trunk** option:

```
interface-mode trunk;
```

You can include this statement at the following hierarchy levels:

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

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Configuring the VLAN ID List for a Trunk Interface

To configure the list of VLAN IDs to be accepted by the trunk port, include the **vlan-id-list** statement and specify the list of VLAN IDs. You can specify individual VLAN IDs with a space separating the ID numbers, specify a range of VLAN IDs with a dash separating the ID numbers, or specify a combination of individual VLAN IDs and a range of VLAN IDs.

vlan-id-list [*number number-number*];

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]**

When a packet is received that is tagged with a VLAN ID specified in the trunk interface list of VLAN IDs, the packet is accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.

When a packet is received that is tagged with a VLAN ID not specified in the trunk interface list of VLAN IDs, the native VLAN ID is pushed in front of the existing VLAN tag or tags and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

When an untagged packet is received on a trunk interface, the native VLAN ID is added to the packet and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

A bridge domain configured with a matching VLAN ID must be configured before the trunk interface is configured. To learn more about configuring bridge domains, see the *Junos Routing Protocols Configuration Guide*.

Related Documentation

- 802.1Q VLANs Overview on page 53
- *Ethernet Interfaces*

Configuring a Trunk Interface on a Bridge Network

On MX Series routers, you can configure a trunk interface on a bridge network.

The following output sample shows trunk port configuration on a bridge network:

```
user@host# run show interfaces
ge-0/0/0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
}
ge-2/0/0 {
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 1-200;
        }
    }
}
ge-2/0/1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1;
    }
}
```

If you want **igmp-snooping** to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge domain. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed as shown after the sample configuration below:

```
user@host# run show configuration
interfaces {
    ge-5/1/1 {
        flexible-vlan-tagging;
        native-vlan-id 1;
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list 401;
            }
        }
    }
    irb {
        unit 401 {
            family inet {
                address 192.168.2.2/27;
            }
        }
    }
}
protocols {
    igmp {
        interface all;
    }
}
bridge-domains {
    VLAN-401 {
        vlan-id 401;
        routing-interface irb.401;
        protocols {
            igmp-snooping;
        }
    }
}

user@host# commit
[edit bridge-domains]
  'VLAN-401'
```

 Copyright © 2013, Juniper Networks, Inc.

```
    IGMP Snooping not supported with IRB and trunk mode interface ge-5/1/1.0
commit complete
```

To achieve IGMP snooping for a bridge domain, you should use such a configuration as shown in the following example:

```
user@host# run show configuration
interfaces {
    ge-0/0/1 {
        flexible-vlan-tagging;
        native-vlan-id 1;
        encapsulation flexible-ethernet-services;
        unit 0 {
            encapsulation vlan-bridge;
            vlan-id 401;
        }
    }
    irb {
        unit 401 {
            family inet {
                address 192.168.2.2/27;
            }
        }
    }
}
protocols {
    igmp {
        interface all;
    }
}
bridge-domains {
    VLAN-401 {
        vlan-id 401;
        interface ge-0/0/1.0;
        routing-interface irb.401;
        protocols {
            igmp-snooping;
        }
    }
}
```

```
user@host# commit
commit complete
```

Related Documentation

- 802.1Q VLANs Overview on page 53
- *interface-mode*
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 4

Configuring Aggregated Ethernet Interfaces

- Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection on page 189
- Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers on page 190

Aggregated Ethernet Interfaces Overview

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-flow load balancing.

NOTE: For information about configuring circuit cross-connects over aggregated Ethernet, see *Circuit and Translational Cross-Connects Overview*.

Platform Support for Aggregated Ethernet Interfaces

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device and then associating a set of ports that have the same speed and are in full-duplex mode. The physical interfaces can be Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, or 10-Gigabit Ethernet IQ2 and IQ2-E. Generally, you cannot use a combination of these interfaces within the same aggregated link; however, you can combine Gigabit Ethernet and Gigabit Ethernet IQ interfaces in a single aggregated Ethernet bundle.

Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

 Copyright © 2013, Juniper Networks, Inc.

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

TIP:

Going forward:

- Aggregated Ethernet link with member links of different modes will be referred as *10-Gigabit Ethernet mixed mode aggregated Ethernet link*.
- Aggregated Ethernet link with member links of different rates will be referred as *mixed rate aggregated Ethernet link*.
- These aggregated Ethernet links will generically be referred as *mixed aggregated Ethernet links*.

Table 7 on page 83 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 7: Platform Support Matrix for Mixed Aggregated Ethernet Bundles

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
10-Gigabit Ethernet LAN and WAN (WAN rate: OC192)	T640, T1600, T4000, and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (PF-24XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
		• Enhanced Scaling FPC3 (T640-FPC3-ES)	• 10-Gigabit Ethernet PIC with XENPAK (PC-1XGE-XENPAK)
		• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP) • 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)

Copyright © 2013, Juniper Networks, Inc.

Table 7: Platform Support Matrix for Mixed Aggregated Ethernet Bundles *(continued)*

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
40-Gigabit Ethernet, 100-Gigabit Ethernet	T4000 and TX Matrix Plus routers	• T4000 FPC5 (T4000-FPC5-3D)	• 100-Gigabit Ethernet PIC with CFP (PF-1CGE-CFP)
	T640, T1600, T4000, and TX Matrix Plus routers	• Enhanced Scaling FPC4 (T640-FPC4-ES) • Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) • T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)	• 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4) NOTE: This PIC is available packaged only in an assembly with the T1600-FPC4-ES FPC. • 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP)

The following routers support a maximum of 16 physical interfaces per single aggregated Ethernet bundle:

- M120
- M320
- All MX Series 3D Universal Edge Routers
- All T Series routers

All other routers support a maximum of eight physical interfaces per aggregated Ethernet bundle.

On M Series and T Series routers, you can create a maximum of 1024 logical interfaces on an aggregated Ethernet interface.

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, PICs, or MPCs.

Configuration Guidelines for Aggregated Ethernet Interfaces

- Simple filters are not supported for interfaces in aggregated Ethernet bundles:
  - On M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
  - On MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

  For more information about simple filters, see the *Junos OS Class of Service Library for Routing Devices*.
- On the aggregated Ethernet bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific capabilities, see "Gigabit Ethernet Accounting and Policing Overview" on page 395.

 Copyright © 2013, Juniper Networks, Inc.

- Aggregated Ethernet interfaces can be either tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on MX Series routers support the configuration of **flexible-vlan-tagging** and **native-vlan-id** on dual-tagged frames, which consist of the following configuration statements:
  - inner-tag-protocol-id
  - inner-vlan-id
  - pop-pop
  - pop-swap
  - push-push
  - swap-push
  - swap-swap

  In all cases, you must set the number of aggregated Ethernet interfaces on the chassis. You can also set the link speed and the minimum links in a bundle.
- When configuring mixed aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers, consider the following:
  - A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
  - Link Aggregation Control Protocol (LACP), aggregated Ethernet link protection, and LACP link protection are supported only on mixed aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
  - Traffic distribution is based on the hash calculated on the egress packet header. Hash range is fairly distributed according to member links' speed. This guarantees hash fairness but it does not guarantee fair traffic distribution depending on the rate of the egress streams.
  - Packets are dropped when the total throughput of the hash flow exiting a member link (or multiple hash flows exiting a single member link) exceeds the link speed of the member link. This can happen when egress member link changes because of a link failure and the hash flow switches to a member link of speed that is less than the total throughput of the hash flow.
  - Rate-based CoS components such as scheduler, shaper, and policer are not supported on mixed rate aggregated Ethernet links. However, the default CoS settings are supported by default on the mixed rate aggregated Ethernet links.
  - Load balancing is performed at the ingress Packet Forwarding Engine. Therefore, you must ensure that the egress traffic on the aggregated Ethernet link enters through the hardware platforms that support mixed aggregated Ethernet bundles. Table 7 on page 83 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.
  - Mixed aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed aggregated Ethernet load balancing is configured at egress.

Copyright © 2013, Juniper Networks, Inc.

   - Load balancing of the egress traffic across the member links of a mixed rate aggregated Ethernet link is proportional to the rates of the member links.
   - Egress multicast load balancing is not supported on mixed aggregated Ethernet interfaces.
   - Changing the **edit interfaces ae*x* aggregated-ether-options link-speed** configuration of a mixed aggregated Ethernet link, which is configured on the supported interfaces of on T640, T1600, T4000, and TX Matrix Plus routers, leads to aggregated Ethernet link flapping.
   - When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4), ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
   - When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC with CFP, changing aggregated Ethernet link protection or LACP link protection configurations leads to aggregated Ethernet link flapping.
   - For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC with CFP, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP.
   - The **show interfaces ae*x*** command displays the link speed of the aggregated Ethernet interface, which is the sum of the link speeds of all the active member links.

- Use the **show interfaces *aggregate-interface* extensive** and **show interfaces *aggregate.logical-interface*** commands to show the bandwidth of the aggregate. Also, the SNMP object identifier **ifSpeed/ifHighSpeed** shows the corresponding bandwidth on the aggregate logical interface if it is configured properly.

Related Documentation

- inner-tag-protocol-id on page 581
- inner-vlan-id on page 582
- pop-pop on page 623
- pop-swap on page 624
- push-push on page 636
- swap-push on page 654
- swap-swap on page 655
- Configuring Mixed Aggregated Ethernet Links on page 91
- Gigabit Ethernet Accounting and Policing Overview on page 395
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Configuring an Aggregated Ethernet Interface

On Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on M Series and T Series routers, you can associate a physical interface with an aggregated Ethernet interface.

To configure an aggregated Ethernet interface:

1. Specify that you want to configure the link aggregation group interface.

   user@host# **edit interfaces** ***interface-name***

2. Configure the aggregated Ethernet interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x***

You specify the interface instance number ***x*** to complete the link association; ***x*** can be from 0 through 127, for a total of 128 aggregated interfaces on M Series and T Series routers and can be from 1 through 480, for a total of 480 aggregated interfaces on MX Series routers. You must also include a statement defining **ae*x*** at the **[edit interfaces]** hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see "Ethernet Interfaces Overview" on page 35, and for a sample configuration, see "Example: Configuring Aggregated Ethernet Interfaces" on page 544.

NOTE: In general, aggregated Ethernet bundles support the features available on all supported interfaces that can become a member link within the bundle. As an exception, Gigabit Ethernet IQ features and some newer Gigabit Ethernet features are not supported in aggregated Ethernet bundles.

Gigabit Ethernet IQ and SFP interfaces can be member links, but IQ- and SFP-specific features are not supported on the aggregated Ethernet bundle even if all the member links individually support those features.

You need to configure the correct link speed for the aggregated Ethernet interface to eliminate any warning message.

NOTE: Before you commit an aggregated Ethernet configuration, ensure that link mode is not configured on any member interface of the aggregated Ethernet bundle; otherwise, the configuration commit check fails.

Related Documentation

- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 176
- Deleting an Aggregated Ethernet Interface on page 93
- Aggregated Ethernet Interfaces Overview on page 82

- *Ethernet Interfaces*

Configuring Junos OS for Supporting Aggregated Devices

Junos OS supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard.

Tasks for configuring aggregated devices are:

Configuring Virtual Links for Aggregated Devices

To define virtual links, you need to specify the associations between physical and logical devices within the **[edit interfaces]** hierarchy, and assign the correct number of logical devices by including the **device-count** statement at the **[edit chassis aggregated-devices ethernet]** and **[edit chassis aggregated-devices sonet]** hierarchy levels:

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count number;
  }
  sonet {
    device-count number;
  }
}
```

The maximum number of Ethernet logical interfaces that you can configure is 128. On M Series and T Series routers, you can configure a maximum number of 128 aggregated interfaces. On MX Series routers, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces are numbered from **ae0** through **ae127** for M Series and T Series routers, and the aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers. The maximum number of SONET/SDH logical interfaces is 64. The aggregated SONET/SDH interfaces are numbered from **as0** through **as63**.

NOTE: **Starting with Junos OS Release 13.2, a maximum of 64 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces. In releases before Junos OS Release 13.2, a maximum of 16 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces.**

Configuring LACP Link Protection at the Chassis Level

Link Aggregation Control Protocol (LACP) is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and

 Copyright © 2013, Juniper Networks, Inc.

untagged aggregated Ethernet with or without LACP enabled. LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP link protection enables you to force active and standby links within an aggregated Ethernet. You configure LACP link protection by using the **link-protection** and **system-priority** statements at either the chassis or interface level and by configuring port priority at the interface level using the **system-priority** statement. Configuring LACP parameters at the chassis level results in all aggregated Ethernet interfaces using the defined values unless overridden by the LACP configuration on a specific interface.

```
[edit chassis]
aggregated-devices {
   ethernet {
      lacp {
         link-protection {
            non-revertive;
         }
         system-priority priority;
      }
   }
}
```

NOTE: **LACP link protection also uses port priority. You can configure port priority at the Ethernet interface [gigether-options] hierarchy level using the port-priority statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).**

Enabling LACP Link Protection

To enable LACP link protection for aggregated Ethernet interfaces on the chassis, use the **link-protection** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
link-protection {
   non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, after a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch, and the current link remains active.

Copyright © 2013, Juniper Networks, Inc.

CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

Configuring System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the chassis, use the **system-priority** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
system-priority priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 through 65,535.

Configuring the Maximum Links Limit

To configure the maximum links limit, use the **maximum-links** statement at the **[edit chassis aggregated-devices]** hierarchy level:

```
[edit chassis aggregated-devices]
maximum-links maximum-links-limit;
```

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 87
- *Ethernet Interfaces*
- Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers on page 481
- *Configuring Aggregated SONET/SDH Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Configuring Mixed Aggregated Ethernet Links

In releases before Junos OS Release 13.2, all interfaces that form an aggregated Ethernet bundle must have the same speed and must be in full-duplex mode. Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of mixed rate or mixed mode aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers, you need to configure the **mixed** option for the [**edit interfaces ae*x* aggregated-ether-options link-speed**] statement.

To configure mixed aggregated Ethernet interfaces:

1. Configure the number of aggregated logical devices available to the router:

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count number
   ```

   For example:

   ```
   [edit chassis]
   user@host# set aggregated-devices ethernet device-count 3
   ```

2. Configure the minimum number of links that is required for the aggregated Ethernet interface to be labeled *up*:

   ```
   [edit interfaces]
   user@host# set aex aggregated-ether-options minimum-links number
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options minimum-links 2
   ```

NOTE: By default, only one link needs to be up for the bundle to be labeled *up*.

3. Configure the **link-speed** statement. Specify the **mixed** option for the **link-speed** statement to indicate the mixed aggregated Ethernet bundle configuration.

   ```
   [edit interfaces]
   user@host# set aex aggregated-ether-options link-speed mixed
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae0 aggregated-ether-options link-speed mixed
   ```

   NOTE: It is mandatory to configure the **mixed** option when configuring the mixed aggregated Ethernet bundle on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4). On other supported platforms, if the **[edit interfaces ae*x* aggregated-ether-options link-speed]** statement is not configured, the mixed configuration is applied by default.

4. Configure the members links of the aggregated Ethernet bundle:

   ```
   [edit interfaces]
   user@host# set interface-name gigether-options 802.3ad aex
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set xe-0/0/1 gigether-options 802.3ad ae0
   user@host# set et-1/1/0 gigether-options 802.3ad ae0
   user@host# set ce-1/1/1 gigether-options 802.3ad ae0
   ```

5. Configure an interface family for the aggregated Ethernet bundle:

   ```
   [edit interfaces]
   user@host# set aex unit number family inet address address
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae0 unit 0 family inet address 100.100.100.1/30
   ```

6. Commit the configuration:

   ```
   [edit]
   user@host# commit
   ```

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 82
- Configuring Aggregated Ethernet Link Speed on page 186
- link-speed on page 596

Copyright © 2013, Juniper Networks, Inc.

Deleting an Aggregated Ethernet Interface

There are two approaches to deleting an aggregated Ethernet interface:

- You can delete an aggregated Ethernet interface from the interface configuration. The Junos OS removes the configuration statements related to **ae***x* and sets this interface to down state.
- You can also permanently remove the aggregated Ethernet interface from the device configuration by deleting it from the device-count on the routing device.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

   This step changes the interface state to down and removing the configuration statements related to **ae***x*.

   [edit]
   user@host# **delete interfaces ae***x*

2. Delete the interface from the device count.

   [edit]
   user@host# **delete chassis aggregated-devices ethernet device-count**

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 87
- Configuring the Number of Aggregated Ethernet Interfaces on the Device on page 176
- Aggregated Ethernet Interfaces Overview on page 82

- *Ethernet Interfaces*

Configuring Multichassis Link Aggregation

On MX Series routers and EX Series switches, multichassis link aggregation (MC-LAG) enables a device to form a logical LAG interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node level redundancy, multi-homing support, and loop-free Layer 2 network without running Spanning Tree Protocol (STP). MC-LAG can be configured for VPLS routing instance, CCC application, and Layer 2 circuit encapsulation types.

The MC-LAG devices use Inter-Chassis Communication Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

On one end of MC-LAG is a MC-LAG client device that has one or more physical links in a link aggregation group (LAG). This client device does not need to be aware of MC-LAG. On the other side of MC-LAG are two MC-LAG network devices. Each of these network devices has one or more physical links connected to a single client device. The network devices coordinate with each other to ensure that data traffic is forwarded properly.

MC-LAG includes the following functionality:

- Active standby mode is supported using Link Aggregation Control Protocol (LACP)
- MC-LAG operates only between two chassis.
- Layer 2 circuit functions are supported with **ether-ccc** encapsulation.
- VPLS functions are supported with **ether-vpls** and **vlan-vpls**.

NOTE: Ethernet connectivity fault management (CFM) specified in IEEE 802.1ag standard for Operation, Administration, and Management (OAM) is *not* supported on MC-LAG interfaces.

To enable MC-LAG, include the **mc-ae** statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level along with either the **ethernet-bridge**, **encapsulation ethernet-ccc**, **encapsulation ethernet-vpls**, or **flexible-ethernet-services** statement at the **[edit interfaces ae*X*]** hierarchy level. You also need to configure the **lacp** statement and the **admin-key** and **system-id** statements at the **[edit interfaces ae*X* aggregated-ether-options] hierarchy level**:

```
[edit interfaces aeX]
encapsulation (ethernet-bridge | ethernet-ccc | ethernet-vpls | flexible-ethernet-services);
aggregated-ether-options {
  lacp {
    active;
    admin-key number;
    system-id mac-address;
    system-priority number;
  }
  mc-ae {
    chassis-id chassis-id;
    events {
      iccp-peer-down {
        force-icl-down;
        prefer-status-control-active;
      }
    }
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    status-control (active | standby);
  }
}
```

NOTE: When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning.

 Copyright © 2013, Juniper Networks, Inc.

To delete a MC-LAG interface from the configuration, issue the **delete interfaces ae*X* aggregated-ether-options mc-ae** command at the **[edit]** hierarchy level in configuration mode:

```
[edit]
user@host# delete interfaces aeX aggregated-ether-options mc-ae
```

To prevent loops in MC-LAG topologies, configure the two edge nodes with same (STP) virtual root ID using Reverse L2 Gateway Protocol (RL2GP). This root ID should be superior to all bridges in the downstream network while downstream bridges have to be capable of running STP. RL2GP should be configured on both MC-LAG nodes to prevent loops. A potential loop, such as one that can happen due to improper cabling at the core or the access switching layer, or due to a bug in server software, is broken by STP blocking one of the interfaces in the downstream network. Because both the MC-LAG nodes are root bridges (virtual), the MC-LAG interface remains in the forwarding state. Downstream bridge receives bridge protocol data units (BPDUs) from both the nodes and thus receives twice the number of BPDUs on its **ae** interface. If both MC-LAG nodes use the same **ae** interface name, STP port number will be identical, which reduces the STP load on downstream bridge.

Perform the following steps on each switch that is hosting an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-optionsmc-ae mc-ae-id mc-ae-id
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-optionsmc-ae mc-ae-id 3
   ```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-optionsmc-ae chassis-id chassis-id
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-optionsmc-ae chassis-id 0
   ```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.

   NOTE: Only active-active mode is supported at this time.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-optionsmc-ae mode mode
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-optionsmc-ae mode active-active
   ```

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary. Primary is **active**, and secondary is **standby**.

NOTE: You must configure status control on both switches hosting the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
user@host# set aeX aggregated-ether-optionsmc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
user@host# set aeX aggregated-ether-optionsmc-ae status-control active
```

5. Specify the same LACP system ID on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-optionslacp system-id mac-address
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
   ```

6. Specify the same LACP administration key on each switch.

   ```
   [edit interfaces]
   user@host# set aeX aggregated-ether-optionslacp admin-key number
   ```

   For example:

   ```
   [edit interfaces]
   user@host# set ae1 aggregated-ether-optionslacp admin-key 3
   ```

7. Configure ICCP by doing the following on each switch hosting the MC-LAG:

   a. Configure the local IP address to be used by all switches hosting the MC-LAG.

      ```
      [edit protocols]
      user@host# set iccp local-ip-addr local-ip-address
      ```

      For example:

      ```
      [edit protocols]
      user@host# set iccp local-ip-addr 3.3.3.1
      ```

   b. (Optional) Configure the IP address of the router and the time during which an ICCP connection must succeed between the routeres hosting the MC-LAG.

      Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

      ```
      [edit protocols]
      user@host# set iccp peer peer-ip-address session-establishment-hold-time seconds
      ```

      For example:

      ```
      [edit protocols]
      user@host# set iccp peer 3.3.3.2 session-establishment-hold-time 50
      ```

   c. (Optional) Configure the IP address to be used for backup liveness detection:

      NOTE: By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

 Copyright © 2013, Juniper Networks, Inc.

[edit protocols]
user@host# **set iccp peer *peer-ip-address* backup-liveness-detection backup-peer-ip *ip-address***

For example:

[edit protocols]
user@host# **set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232**

d. Configure the minimum interval at which the router must receive a reply from the other router with which it has established a Bidirectional Forwarding Detection (BFD) session.

NOTE: Configuring the minimum receive interval is required to enable BFD.

[edit protocols]
user@host# **set iccp peer *peer-ip-address* liveness-detection minimum-receive-interval *seconds***

For example:

[edit protocols]
user@host# **set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 60**

e. Configure the minimum transmit interval during which a router must receive a reply from a router with which it has established a BFD session.

[edit protocols]
user@host# **set iccp peer *peer-ip-address* liveness-detection transmit-interval minimum-interval *seconds***

For example:

[edit protocols]
user@host# **set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 60**

8. Configure a multichassis protection link between the routeres.

[edit]
user@host# **set multi-chassis multi-chassis-protection *peer-ip-address* interface *interface-name***

For example:

[edit protocols]
user@host# **set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0**

9. Enable RSTP globally on all interfaces.

[edit]
user@host# **set protocols rstp interface all mode point-to-point**

10. Disable RSTP on the ICL-PL interfaces on both routeres.

[edit]
user@host# **set protocols rstp interface *interface-name*** disable

For example:

[edit]
user@host# **set protocols rstp interface ae0.0** disable

11. Configure the MC-LAG interfaces as edge ports on both routeres.

set protocols rstp interface *interface-name*

For example:

```
[edit]
user@host# set protocols rstp interface ae1.0
```

12. Enable BPDU block on all interfaces except for the ICL-PL interfaces on both routeres.

```
[edit]
user@host# set protocols rstp bpdu-block-on-edge
```

For example:

```
[edit]
user@host# set protocols rstp bpdu-block-on-edge
```

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 98
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 110
- *show interfaces mc-ae*
- *Ethernet Interfaces*

Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview

MX Series routers support active-active bridging and virtual router redundancy protocol (VRRP) over Integrated routing and bridging (IRB). This is a common scenario used in data centers. This section provides an overview of the supported functionality.

Active-active bridging and VRRP over IRB support extends multichassis link aggregation group (MC-LAG) by adding the following functionality:

- Interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for active-active bridging
- Active-active bridging
- VRRP over IRB for active-active bridging
- A single bridge domain cannot correspond to two redundancy group IDs

The following functionalities are supported for MC-LAG in an active-active bridging domain:

- MC-LAG is supported only between two chassis. Interchassis link (ICL) pseudowire interface or Ethernet interface (ICL-PL field) for active-active bridging Active-active bridging VRRP over IRB for active-active bridging.
- For VPLS networks, you can configure the aggregated Ethernet (aeX) interfaces on MC-LAG devices with the **encapsulation ethernet-vpls** statement to use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values or the **encapsulation vlan-vpls** statement to use Ethernet VLAN encapsulation on VPLS circuits.
- Layer 2 circuit functionalities are supported with **ethernet-ccc** as the encapsulation mode.

Copyright © 2013, Juniper Networks, Inc.

- Network topologies in a triangular and square pattern are supported. In a triangular network design, with equal-cost paths to all redundant nodes, slower, timer-based convergence can possibly be prevented. Instead of indirect neighbor or route loss detection using hellos and dead timers, you can identify the physical link loss and denote a path as unusable and reroute all traffic to the alternate equal-cost path. In a square network design, depending on the location of the failure, the routing protocol might converge to identify a new path to the subnet or the VLAN, causing the convergence of the network to be slower.
- Interoperation of Link Aggregation Control Protocol (LACP) for MC-LAG devices. LACP is one method of bundling several physical interfaces to form one logical interface. When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as active for the link to be up.
- Active-standby mode is supported using LACP. When an MC-LAG operates in the active-standby mode, one of the router's ports only becomes active when failure is detected in the active links. In this mode, the PE routers perform an election to determine the active and standby routers.
- Configuration of the pseudowire status type length variable (TLV) is supported. The pseudowire status TLV is used to communicate the status of a pseudowire back and forth between two provider edge (PE) routers. The pseudowire status negotiation process assures that a PE router reverts back to the label withdraw method for pseudowire status if its remote PE router neighbor does not support the pseudowire status TLV.
- The MC-LAG devices use Interchassis Communication Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

Keep the following points in mind when you configure MC-LAG in an active-active bridging domain:

- A single bridge domain cannot be associated with two redundancy groups. You cannot configure a bridge domain to contain logical interfaces from two different multichassis aggregated Ethernet (MC-AE) interfaces and associate them with different redundancy group IDs by using the redundancy group group-id statement at the **[edit interfaces ae*X* aggregated-ether-options]** hierarchy level.
- You must configure logical interfaces in a bridge domain from a single MC-AE interface and associate it with a redundancy group. You must configure a service ID by including the **service-id *vid*** statement at the **[edit bridge-domains *bd-name*]** hierarchy level for MC-AE interfaces if you configure logical interfaces on MC-AE interfaces that are part of the bridge domain.

For a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain and with VRRP configured over an integrated routing and bridging (IRB) interface, you must include the **accept-data** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other layer 3 protocols and applications to work properly over multi-chassis aggregated Ethernet (mc-aeX) interfaces.

NOTE: On an MC-LAG, the unit number associated with aggregated Ethernet interfaces on provider Edge Router, PE1, must match the unit number associated with aggregated Ethernet interfaces on provider Edge Router, PE2. If the unit numbers differ, MAC address synchronization does not happen. As a result, the status of the MAC address on the remote provider edge router remains in pending state.

In an IPv6 network, you cannot configure an MC-LAG in an active-active bridge domain if you specified the **vlan-id none** statement at **[edit bridge-domain *bd-name*]** hierarchy level. The **vlan-id none** statement that enables the removal of the incoming VLAN tags identifying a Layer 2 logical interface when packets are sent over VPLS pseudowires is not supported for IPv6 packets in an MC-LAG.

The topologies shown in Figure 1 on page 100 and Figure 2 on page 101 are supported. These figures use the following abbreviations:

- Aggregated Ethernet (AE)
- Interchassis link (ICL)
- Multichassis link (MCL)

Figure 1: Single Multichassis Link

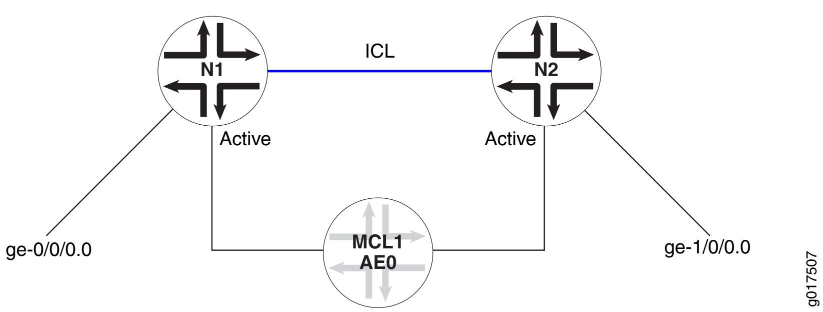

 Copyright © 2013, Juniper Networks, Inc.

Figure 2: Dual Multichassis Link

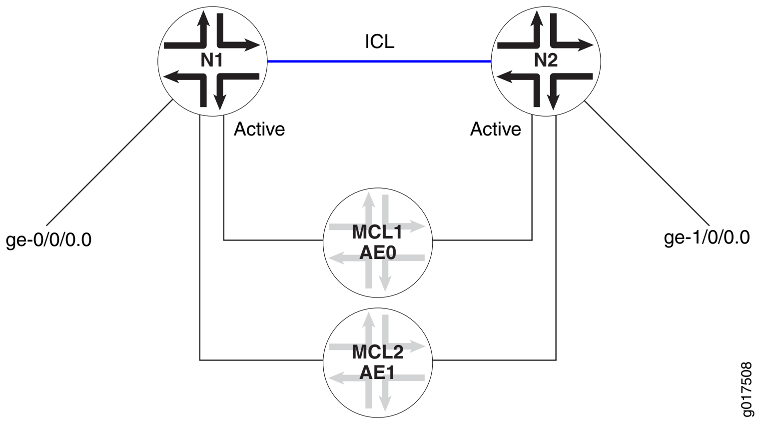

The following functionality is not supported:

- Virtual private LAN service (VPLS) within the core
- Bridged core
- Topology as described in Rule 4 of “Data Traffic Forwarding Rules” on page 103
- Routed multichassis aggregated Ethernet (RMC-AE), where the VRRP backup master is used in the edge of the network
- Track object, where in the case of an MC-LAG, the status of the uplinks from the provider edge can be monitored and the MC-LAG can act on the status
- Mixed mode (active-active MC-LAG is supported on MX series routers with MPC/MIC interfaces only). All interfaces in the bridge-domain that are mc-ae active-active, must be on MPC/MICs.

The topologies shown in Figure 3 on page 101, Figure 4 on page 102 and Figure 5 on page 102 are not supported:

Figure 3: Interchassis Data Link Between Active-Active Nodes

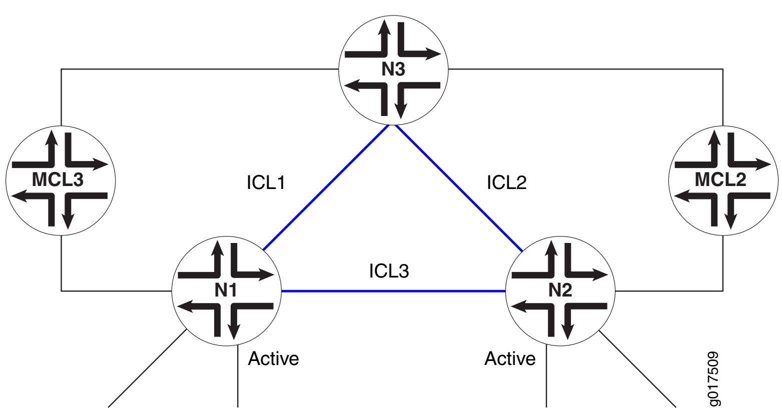

Figure 4: Active-Active MC-LAG with Single MC-LAG

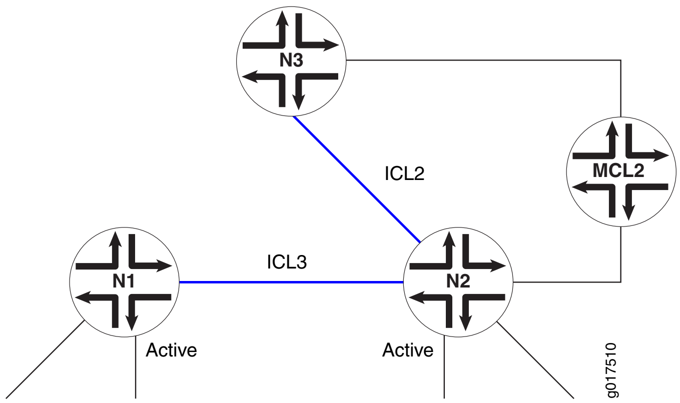

Figure 5: Active-Active MC-LAG with Multiple Nodes on a Single Multichassis Link

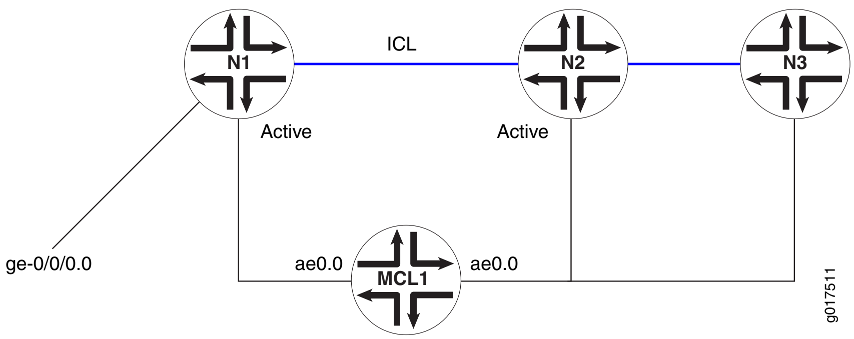

NOTE: A redundancy group cannot span more than two routers.

When configured to be active-active, the client device load balances the traffic to the peering MC-LAG network devices. In a bridging environment, this could potentially cause the following problems:

- Traffic received on the MC-LAG from one MC-LAG network device could be looped back to the same MC-LAG on the other MC-LAG network device.
- Duplicated packets could be received by the MC-LAG client device.
- Traffic could be unnecessarily forwarded on the interchassis link.

To better illustrate the problems listed above, consider Figure 6 on page 103, where an MC-LAG device MCL1 and single-homed clients **ge-0/0/0.0** and **ge-1/0/0.0** are allowed to talk to each other through an ICL:

- Traffic received on network routing instance N1 from MCL1 could be flooded to ICL to reach network routing instance N2. Once it reaches network routing instance N2, it could be flooded back to MCL1.

 Copyright © 2013, Juniper Networks, Inc.

- Traffic received on interface **ge-0/0/0.0** could be flooded to MCL1 and ICL on network routing instance N1. Once network routing instance N2 receives such traffic from ICL, it could be again flooded to MCL1.
- If interface **ge-1/0/0.0** does not exist on network routing instance N2, traffic received from interface **ge-0/0/0.0** or MCL1 on network routing instance N1 could be flooded to network routing instance N2 through ICL unnecessarily since interface **ge-0/0/0.0** and MCL1 could reach each other through network routing instance N1.

Figure 6: MC-LAG Device and Single-Homed Client

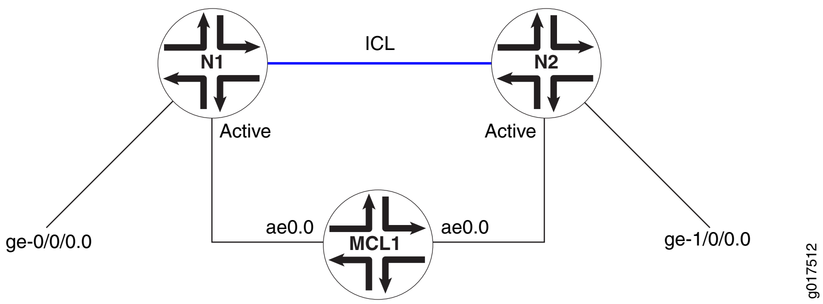

Advantages of Using Multichassis Link Aggregation Groups

An MC-LAG reduces operational expenses by providing active-active links with a LAG, eliminates the need for the Spanning Tree Protocol (STP), and provides faster layer 2 convergence upon link and device failures.

An MC-LAG adds node-level redundancy to the normal link-level redundancy that a LAG provides. An MC-LAG improves network resiliency, which reduces network down time as well as expenses.

In data centers, it is desirable for servers to have redundant connections to the network. You probably want active-active connections along with links from any server to at least two separate switches.

An MC-LAG allows you to bond two or more physical links into a logical link between two switches or between a server and a switch, which improves network efficiency. An MC-LAG enables you to load balance traffic on multiple physical links. If a link fails, the traffic can be forwarded through the other available link and the logical aggregated link remains in the UP state.

Data Traffic Forwarding Rules

In active-active bridging and VRRP over IRB topographies, network interfaces are categorized into three different interface types, as follows:

S-Links—Single-homed link (S-Link) terminating on MC-LAG-N device or MC-LAG in active-standby mode. In Figure 6 on page 103, interfaces **ge-0/0/0.0** and **ge-1/0/0.0** are S-Links.

MC-Links—MC-LAG links. In Figure 6 on page 103, interface **ae0.0** is the MC-Link.

ICL—Interchassis data link.

Based on incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations, as described in the data traffic forwarding rules. Note that if only one of the MC-LAG member link is in the UP state, it is considered an S-Link.

The following data traffic forwarding rules apply:

1. When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.

2. NOTE: **The functionality described in rule 2 is not supported.**

   The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:

   a. If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device.

   b. Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other.

   Only if both a. and b. are true, is traffic always forwarded to the ICL.

3. When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network that the packet comes from.

4. NOTE: **The topology shown in Figure 7 on page 105 is not supported.**

   In certain cases, for example the topology shown in Figure 7 on page 105, there could be a loop caused by the ICL. To break the loops, one of the following mechanisms could be used:

   a. Run certain protocols, such as spanning tree protocol (STP). In this case, whether packets received on one ICL are forwarded to other ICLs is determined by using Rule 3.

   b. Configure the ICL to be fully meshed among the MC-LAG network devices. In this case, traffic received on the ICL would be not be forwarded to any other ICLs.

   In either case, duplicate packets could be forwarded to the MC-LAG clients. Consider the topology shown in Figure 7 on page 105, where if network routing instance N1 receives a packet from **ge-0/0/0.0**, it could be flooded to ICL1 and ICL3.

   When receiving from ICL1 and ICL3, network routing instances N3 and N2 could flood the same packet to MCL2, as shown in Figure 7 on page 105. To prevent this from happening, the ICL designated forwarder should be elected between MC-LAG peers

 Copyright © 2013, Juniper Networks, Inc.

and traffic received on an ICL could be forwarded to the active-active MC-LAG client by the designated forwarder only.

Figure 7: Loop Caused by the ICL Links

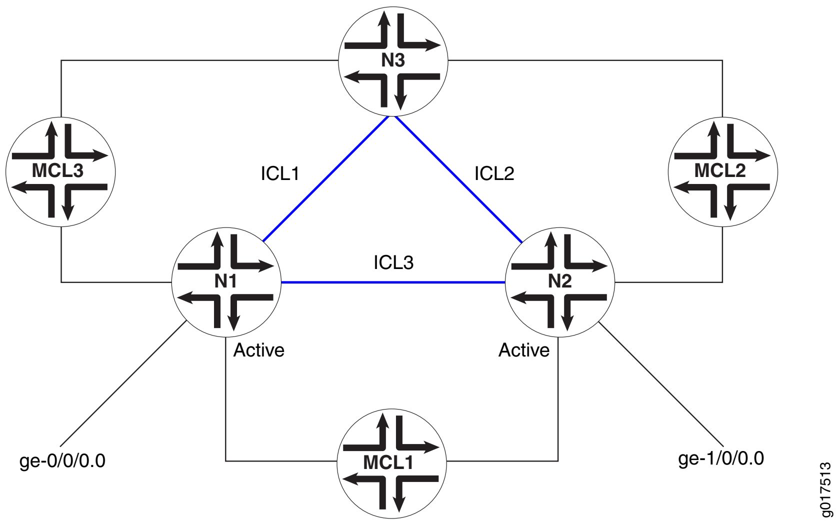

5. When received from an ICL, traffic should not be forwarded to the core-facing client link connection between two provider edge (PE) devices (C-Link) if the peer chassis's (where the traffic is coming from) C-Link is UP.

MAC Address Management

If an MC-LAG is configured to be active-active, upstream and downstream traffic could go through different MC-LAG network devices. Since the media access control address (MAC address) is learned only on one of the MC-LAG network devices, the reverse direction's traffic could be going through the other MC-LAG network and flooded unnecessarily. Also, a single-homed client's MAC address is only learned on the MC-LAG network device it is attached to. If a client attached to the peer MC-LAG network needs to communicate with that single-homed client, then traffic would be flooded on the peer MC-LAG network device. To avoid unnecessary flooding, whenever a MAC address is learned on one of the MC-LAG network devices, it gets replicated to the peer MC-LAG network device. The following conditions should be applied when MAC address replication is performed:

- MAC addresses learned on a MC-LAG of one MC-LAG network device should be replicated as learned on the same MC-LAG of the peer MC-LAG network device.
- MAC addresses learned on single-homed customer edge (CE) clients of one MC-LAG network device should be replicated as learned on ICL-PL interface of the peer MC-LAG network device.
- MAC addresses learned on MC-LAG VE clients of one MC-LAG network device should be replicated as learned on the corresponding VE interface of the peer MC-LAG network device.
- MAC address learning on an ICL is disabled from the data path. It depends on software to install MAC addresses replicated through interchassis control protocol (ICCP).

Copyright © 2013, Juniper Networks, Inc.

MAC Aging

MAC aging support in the Junos OS extends aggregated Ethernet logic for a specified MC-LAG. A MAC address in software is deleted until all Packet Forwarding Engines have deleted the MAC address. In the case of an MC-LAG, a remote provider edge is treated as a remote Packet Forwarding Engine and has a bit in the MAC data structure.

Layer 3 Routing

In general, when an MC-LAG is configured to provide Layer 3 routing functions to downstream clients, the MC-LAG network peers should be configured to provide the same gateway address to the downstream clients. To the upstream routers, the MC-LAG network peers could be viewed as either equal-cost multi path (ECMP) or two routes with different preference values.

Junos OS supports active-active MC-LAGs by using VRRP over IRB. Junos OS also supports active-active MC-LAGs by using IRB MAC address synchronization. You must configure IRB using the same IP address across MC-LAG peers. IRB MAC synchronization is supported on 32-bit interfaces and interoperates with earlier MPC/MIC releases.

To ensure that Layer 3 operates properly, instead of dropping the Layer 3 packet, the VRRP slave attempts to perform routing functions if the packet is received on an MC-LAG. A VRRP slave sends and responds to address resolution protocol (ARP) requests.

For ARP, the same issue exists as with Layer 2 MAC addresses. Once ARP is learned, it must be replicated to the MC-LAG through ICCP. The peer must install an ARP route based on the ARP information received through ICCP.

For ARP aging, ARP requests on the MC-LAG peers can be aged out independently.

Address Resolution Protocol Active-Active MC-LAG Support Methodology

Suppose one of the PE routers issues an ARP request and another PE router gets the response and, because of the aggregated Ethernet distribution logic, the ARP resolution is not successful. Junos OS uses ARP response packet snooping to perform active-active multichassis link aggregation group support, providing easy synchronization without the need to maintain any specific state.

IGMP Snooping on Active-Active MC-LAG

For multicast to work in an active-active MC-LAG scenario, the typical topology is as shown in Figure 8 on page 107 and Figure 9 on page 108 with interested receivers over S-links and MC-Links. Starting in Junos OS Release 11.2, support is extended for sources connected over Layer 2 interface.

If an MC-LAG is configured to be active-active, reports from MC-LAG clients could reach any of the MC-LAG network device peers. Therefore the IGMP snooping module needs to replicate the states such that the Layer 2 multicast route state on both peers are the same. Additionally for S-Link clients, snooping needs to replicate these joins to its snooping peer, which in the case of Layer 3 connected source, passes this information to the PIM on IRB to enable the designated router to pull traffic for these groups,

Copyright © 2013, Juniper Networks, Inc.

The ICL should be configured as a router facing interface. For the scenario where traffic arrives via a Layer 3 interface, it is a requirement to have PIM and IGMP enabled on the IRB interface configured on the MC-LAG network device peers.

Figure 8: Multicast Topology with Source Connected via Layer 3

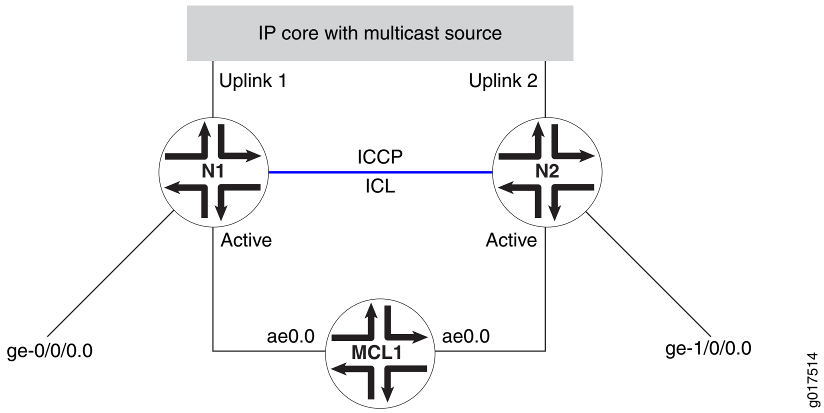

With reference to Figure 8 on page 107, either N1 or N2 becomes a designated router (for this example, N1 is the designated router). Router N1 would therefore pull the multicast traffic from the core. Once multicast data hits the network device N1, the data is forwarded based on the snooping learned route.

For MC-Link clients, data is forwarded via N1. In the case of failover of the MC-Links, the data reaches the client via N2. For S-Link clients on N1, data would be forwarded via normal snooping routes.

For S-Link clients on N2, data is forwarded via the ICL interface. Layer 2 multicast routes on N1 do not show these groups unless there is interest for the same group over MC-Links or over S-Links on N1. For IRB scenario, the IGMP membership and Layer 3 multicast route on N1 does however show these groups learned over the IRB interface.

Therefore, for a case where a specific group interest is only on the S-Link on N2, data arriving on N1 reaches N2 via the default route and the Layer 2 multicast route on N2 has the S-Link in the outgoing interface list.

Copyright © 2013, Juniper Networks, Inc.

Figure 9: Multicast Topology with Source Connected via MC-Link

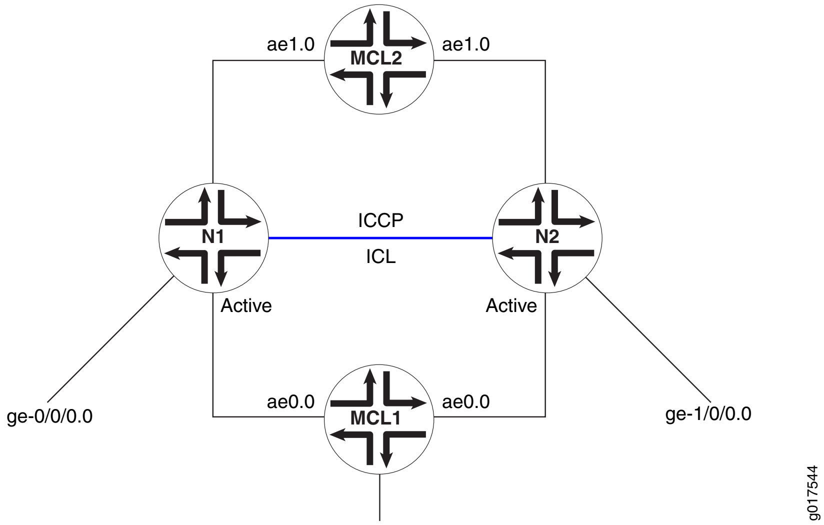

In Figure 9 on page 108, MCL1 and MCL2 are on different devices and the multicast source or IGMP querier is connected via MCL2. The data forwarding behavior seen is similar to that explained for multicast topology with source connected via Layer 3.

NOTE: **IGMP snooping should not be configured in proxy mode. There should be no IGMP hosts or IGMP/PIM routers sitting on the ICL interface.**

Up and Down Event Handling

The following conditions apply to up and down event handling:

1. If the interchassis control protocol (ICCP) connection is UP but the ICL interface becomes DOWN, the router configured as standby will bring down all the multichassis aggregated Ethernet interfaces shared with the peer which is connected to ICL. This will make sure that there are no loops in the network. Otherwise, both PEs will becomes PIM designated routers and, hence, forward multiple copies of the same packet to the customer edge.
2. If the ICCP connection is UP and the ICL comes UP, the router configured as standby will bring up the multichassis aggregated Ethernet interfaces shared with the peer.
3. If both the ICCP connection and the ICL are DOWN, the router configured as standby will bring up the multichassis aggregated Ethernet interfaces shared with the peer.
4. The layer 2 address learn daemon (l2ald) does not store the information about a MAC address learned from a peer in the kernel. If l2ald restarts, and if the MAC address was not learned from the local multichassis aggregated Ethernet interface, l2ald will clear the MAC addresses and this will cause the router to flood the packets destined to this MAC address. This behavior is similar to that in a Routing Engine switchover. (Please note that currently l2ald runs on a Routing Engine only when it is a master). Also, during the time l2ald is DOWN, ARP packets received from an ICCP peer will be

Copyright © 2013, Juniper Networks, Inc.

dropped. ARP retry will take care of this situation. This will be the case with Routing Engine switchover too.

5. If ICCP restarts, l2ald does not identify that a MAC address was learned from a peer and, if the MAC address was learned only from the peer, that MAC address will be deleted and the packets destined to this MAC address will be flooded.

VRRP Active-Standby Support

VRRP in active-standby mode enables Layer 3 routing over the MC-AE interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The virtual routers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address, known as a routed VLAN interface (RVI), maps to either of the VRRP MAC addresses or the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets. At any time, one of the VRRP routers is the master (active), and the other is a backup (standby). Both VRRP active and VRRP backup routers forward Layer 3 traffic arriving on the MC-AE interface. If the master router fails, all the traffic shifts to the MC-AE link on the backup router.

NOTE: **You must configure VRRP on both MC-LAG peers in order for both the active and standby members to accept and route packets. Additionally, configure the VRRP backup router to send and receive ARP requests.**

Routing protocols run on the primary IP address of the RVI, and both of the MC-LAG peers run routing protocols independently. The routing protocols use the primary IP address of the RVI and the RVI MAC address to communicate with the MC-LAG peers. The RVI MAC address of each MC-LAG peer is replicated on the other MC-LAG peer and is installed as a MAC address that has been learned on the ICL-PL.

In many cases, MC-LAG devices are Layer 3 routing devices and perform the default gateway functionality for hosts that are part of the attached IP subnets. The MC-LAG device-pair can therefore share the default gateway routing functionality with the VRRP protocol. The VRRP active-standby operation can be optimized to be an active-active mode of processing because traffic flowing from an MC-LAG client to an MC-LAG device is always sent on one of the available links of an MC-LAG and reaches exactly one of the MC-LAG destination devices. Because of this behavior, both MC-LAG devices in the pair can be enabled as routers for IP traffic destined to the VRRP destination MAC address. Both MC-LAG devices must be full members of the routing domain and have routing entries that allow them to reach the IP destination networks.

The active-active functionality works without changing any VRRP state machine and by activating the routing function in the forwarding plane of the VRRP backup system (similar to the VRRP master). The mechanism enables traffic forwarding in Layer 3 to be fully redundant, leveraging all available link bandwidth. All routers forward traffic and thereby load share routed traffic.

Copyright © 2013, Juniper Networks, Inc.

Interchassis Control Protocol

Interchassis control protocol (ICCP) is used to synchronize configurations, states, and data.

ICCP supports the following types of state information:

- MC-LAG members and their operational states.
- Single-homed members and their operational states.

ICCP supports the following application database synchronization parameters:

- MAC addresses learned and to be aged.
- ARP info learned over IRB.

Interchassis Control Protocol Message

ICCP messages and attribute-value pairs (AVPs) are used for synchronizing MAC address and ARP information.

Related Documentation

- Configuring Multichassis Link Aggregation on page 93
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 110
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 115

- *multi-chassis-protection*
- *peer*
- *show interfaces mc-ae*
- *Ethernet Interfaces*

Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers

The following sections describe the configuration of active-active bridging and VRRP over IRB in multichassis link aggregation (MC-LAG) on MX Series routers:

Copyright © 2013, Juniper Networks, Inc.

Configuring MC-LAG

An MC-LAG is composed of logical link aggregation groups (LAGs) and is configured under the **[edit interfaces ae*X*]** hierarchy, as follows:

```
[edit]
interfaces {
  ae0 {
    encapsulation ethernet-bridge;
    multi-chassis-protection {
      peer 10.10.10.10 {
        interface ge-0/0/0;
        }
      }
    aggregated-ether-options {
      mc-ae {
        mode active-active; # see note below
      }
    }
  }
```

NOTE: **The mode active-active statement is valid only if encapsulation is ethernet-bridge or extended-vlan-bridge.**

Use the **mode** statement to specify if a MC-LAG is **active-standby** or **active-active**. If the ICCP connection is UP and ICL comes UP, the router configured as standby will bring up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer.

Using **multi-chassis-protection** at the physical interface level is a way to reduce the configuration at the logical interface level.

If the following assumption exists (follow the above example):

If there are n+1 logical interfaces under **ae0**, from **ae0.0** through **ae0.n**, there will be n+1 logical interfaces under **ge-0/0/0** as well, from **ge-0/0/0.0** through **ge-0/0/0.n**, each **ge-0/0/0** logical interface will be a protection link for the **ae0** logical interface.

NOTE: **A bridge domain cannot have MC-AE logical interfaces which belong to different redundancy groups.**

Configuring Interchassis Link Label

The interchassis link-protection link (ICL-PL) provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL-PL is an aggregated Ethernet interface. You can configure only one ICL-PL between the two peers, although you can configure multiple MC-LAGs between them.

Copyright © 2013, Juniper Networks, Inc.

Ethernet as interchassis link label (ICL-PL) (assumes interface **ge-0/0/0.0** is used to protect interface **ae0.0** of MC-LAG-1):

```
[edit]
interfaces {
  ae0 {
    ....
    unit 0 {
      multi-chassis-protection {
        peer 10.10.10.10 {
          interface ge-0/0/0.0;
        }
        ....
      }
      ...
    }
  }
}
```

The protection interface can be an Ethernet type interface like **ge**, **xe**, or an aggregated Ethernet (**ae**) interface.

Configuring Multiple Chassis

A top-level hierarchy is used to specify multichassis-related configuration, as follows:

```
[edit]
multi-chassis {
  multi-chassis-protection {
    peer 10.10.10.10 {
      interface ge-0/0/0;
    }
  }
}
```

The above example specifies interface **ge-0/0/0** as the multichassis protection interface for all the multichassis aggregated Ethernet (MC-AE) interfaces which are also part of the peer. This can be overridden by specifying protection at the physical interface level and the logical interface level.

Configuring Service ID

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. You can configure the service IDs under the level of the hierarchies shown in the following examples:

Global Configuration (default logical system)

```
switch-options {
  service-id 10;
}
bridge-domains {
  bd0 {
    service-id 2;
  }
}
routing-instances {
  r1 {
    switch-options {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            service-id 10;
        }
        bridge-domains {
            bd0 {
                service-id 2;
            }
        }
    }
}
```

Logical Systems

```
logical-system {
    ls1 {
        switch-options {
            service-id 10;
        }
    }
}
logical-system {
    ls1 {
        routing-instances {
            r1 {
                switch-options {
                    service-id 10;
                }
            }
        }
    }
}
```

NOTE: Using a service name per bridge domain is not supported.

The bridge level service ID is required to link related bridge domains across peers, and should be configured with the same value. The **service-id** values share the name space across all bridging and routing instances, and across peers. Thus, duplicate values for service IDs are not permitted across these entities.

The service ID at the bridge domain level is mandatory for type non-single VLAN bridge domains. The service ID is optional for bridge domains with a VID defined. If no service ID is defined in the latter case, it is picked up from the service ID configuration for that routing instance.

NOTE: When this default routing instance (or any other routing instance) which contains a bridge domain containing an MC-AE interface is configured, you must configure a global level **switch-options service-id** ***number***, irrespective of whether the contained bridge domains have specific service IDs configured.

In the example shown in Figure 10 on page 114, network routing instances N1 and N2, both for the same service ID, are configured with same service-id in both N1 and N2. Use of a name string instead of a number is not supported.

Copyright © 2013, Juniper Networks, Inc.

Figure 10: N1 and N2 for the Same Service with Same Service ID

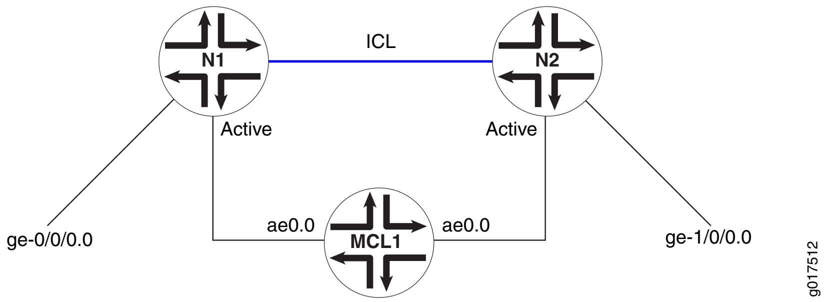

The following configuration restrictions apply:

- The service ID must be configured when the MC-AE interface is configured and an MC-AE logical interface is part of a bridge domain. This requirement is enforced.
- A single bridge domain cannot correspond to two redundancy group IDs.

Figure 11: Bridge Domain with Logical Interfaces from Two MC-AE Interfaces

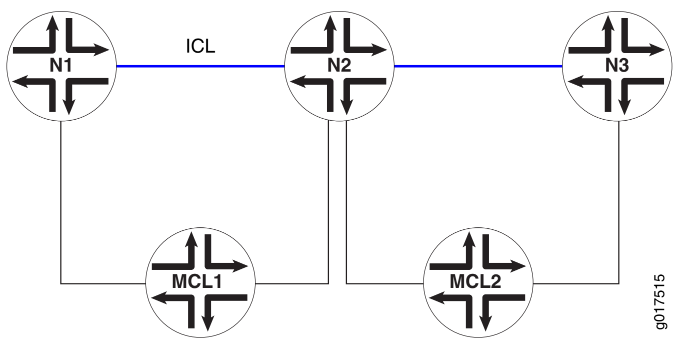

In Figure 11 on page 114, it is possible to configure a bridge domain consisting of logical interfaces from two MC-AE interfaces and map them to a separate redundancy group ID, which is not supported. A service should be mapped one-to-one with the redundancy group providing the service. This requirement is enforced.

To display the MC-AE configuration, use the **show interfaces *mc-ae*** command. For more information, see the *Junos OS Operational Mode Commands*.

Configuring IGMP Snooping for Active-Active MC-LAG

For the multicast solution to work, the following must be configured:

- The multichassis protection link should be configured as a router-facing interface.

```
[edit bridge-domain bd-name]
protocols {
  igmp-snooping {
    interface ge-0/0/0.0 {
      multicast-router-interface;
    }
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

In this example, **ge-0/0/0.0** is an ICL interface.

- The ***multichassis-lag-replicate-state*** statement options should be configured under the ***multicast-snooping-options*** statement for that bridge domain.

NOTE: Snooping with active-active MC-LAG is only supported in non-proxy mode.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 98
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 115
- Configuring Multichassis Link Aggregation on page 93

- *mc-ae*
- *multi-chassis-protection*
- *peer*
- *show interfaces irb*
- *show interfaces mc-ae*
- *Ethernet Interfaces*

Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers

This example illustrates how to configure a multichassis link aggregation group (MC-LAG) in an active-active scenario on MX Series routers.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Overview

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

In Figure 12 on page 117, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facindg the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

Topology Diagram

Figure 12 on page 117 shows the topology used in this example.

Copyright © 2013, Juniper Networks, Inc.

Figure 12: MC-LAG Active-Active on MX Series Routers

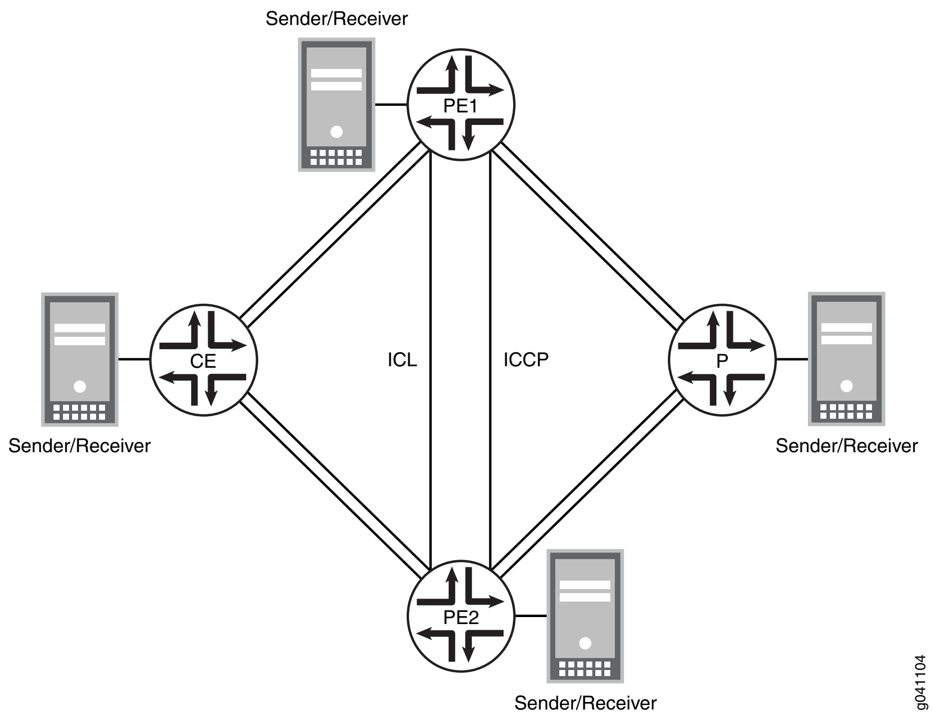

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```

Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
```


Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.2
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@PE1# set aggregated-devices ethernet device-count 5
   ```

2. Specify the members to be included within the aggregated Ethernet bundles.

   ```
   [edit interfaces]
   user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
   user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
   ```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

   ```
   [edit interfaces]
   user@PE1# set ge-1/1/1 flexible-vlan-tagging
   user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
   user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
   user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
   user@PE1# set ge-1/1/4 flexible-vlan-tagging
   user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
   user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
   ```

user@PE1# **set ge-1/1/4 unit 0 vlan-id-range 100-110**
user@PE1# **set ge-1/0/2 unit 0 family inet address 100.100.100.1/30**

4. Configure parameters on the aggregated Ethernet bundles.

   [edit interfaces ae0]
   user@PE1# **set flexible-vlan-tagging**
   user@PE1# **set encapsulation flexible-ethernet-services**
   user@PE1# **set unit 0 encapsulation vlan-bridge**
   user@PE1# **set unit 0 vlan-id-range 100-110**
   user@PE1# **set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0**

   [edit interfaces ae1]
   user@PE1# **set flexible-vlan-tagging**
   user@PE1# **set encapsulation flexible-ethernet-services**
   user@PE1# **set unit 0 encapsulation vlan-bridge**
   user@PE1# **set unit 0 vlan-id-range 100-110**
   user@PE1# **set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0**

5. Configure LACP on the aggregated Ethernet bundles.

   [edit interfaces ae0 aggregated-ether-options]
   user@PE1# **set** lacp **active**
   user@PE1# **set lacp system-priority 100**
   user@PE1# **set lacp system-id 00:00:00:00:00:05**
   user@PE1# **set lacp admin-key 1**

   [edit interfaces ae1 aggregated-ether-options]
   user@PE1# **set lacp active**
   user@PE1# **set lacp system-priority 100**
   user@PE1# **set lacp system-id 00:00:00:00:00:05**
   user@PE1# **set lacp admin-key 1**

6. Configure the MC-LAG interfaces.

   [edit interfaces ae0 aggregated-ether-options]
   user@PE1# **set mc-ae mc-ae-id 5**
   user@PE1# **set mc-ae redundancy-group 10**
   user@PE1# **set mc-ae chassis-id 1**
   user@PE1# **set mc-ae mode active-active**
   user@PE1# **set mc-ae status-control active**

   [edit interfaces ae1 aggregated-ether-options]
   user@PE1# **set mc-ae mc-ae-id 10**
   user@PE1# **set mc-ae redundancy-group 10**
   user@PE1# **set mc-ae chassis-id 1**
   user@PE1# **set mc-ae mode active-active**
   user@PE1# **set mc-ae status-control active**

   The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 118).

Copyright © 2013, Juniper Networks, Inc.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 118) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Copyright © 2013, Juniper Networks, Inc.

```
user@PE1# show bridge-domains
bd0 {
   domain-type bridge;
   vlan-id all;
   service-id 20;
   interface ae1.0;
   interface ge-1/1/1.0;
   interface ge-1/1/4.0;
   interface ae0.0;
}
```

```
user@PE1# show chassis
aggregated-devices {
   ethernet {
      device-count 5;
   }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/6 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-1/0/2 {
   unit 0 {
      family inet {
         address 100.100.100.1/30;
      }
   }
}
ge-1/1/1 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ge-1/1/4 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
```


Copyright © 2013, Juniper Networks, Inc.

```
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 5;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}

user@PE1# show protocols
iccp {
    local-ip-addr 100.100.100.1;
    peer 100.100.100.2 {
        redundancy-group-id-list 10;
        liveness-detection {
            minimum-interval 1000;
        }
```

Copyright © 2013, Juniper Networks, Inc.

```
  }
}
user@PE1# show switch-options
service-id 10;
```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@CE# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
   user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
   ```

 Copyright © 2013, Juniper Networks, Inc.

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-2/1/6]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}
```

```
user@CE# show chassis
```

```
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/0/3 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/1/6 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
      }
   }
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
```

 Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@P# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@P# set ge-1/0/5 gigether-options 802.3ad ae1
   user@P# set ge-1/0/11 gigether-options 802.3ad ae1
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-1/1/3]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-500
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-110
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1 aggregated-ether-options]
   user@P# set lacp active
   user@P# set lacp system-priority 100
   ```

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@P# set vlan-id all
   user@P# set domain-type bridge
   user@P# set interface ge-1/1/3.0
   ```

Copyright © 2013, Juniper Networks, Inc.

```
user@P# set interface ae1.0
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-1/1/3.0;
  interface ae1.0;
}
```

```
user@P# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
user@P# show interfaces
ge-1/0/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/11 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/1/3 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- ***show l2-learning instance*** **extensive**

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 98
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 110
- *Configuring ICCP for MC-LAG*
- *show interfaces (Aggregated Ethernet)* in the *Junos OS Operational Mode Commands*

Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the routers participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Copyright © 2013, Juniper Networks, Inc.

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a router.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a router.
- Configure Virtual Router Redundancy Protocol (VRRP) on a router.

Overview

In this example, you configure an MC-LAG across two routers by including interfaces from both routers in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each router in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In Figure 12 on page 117, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facindg the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

Topology Diagram

Figure 12 on page 117 shows the topology used in this example.

Copyright © 2013, Juniper Networks, Inc.

Figure 13: MC-LAG Active-Active on MX Series Routers

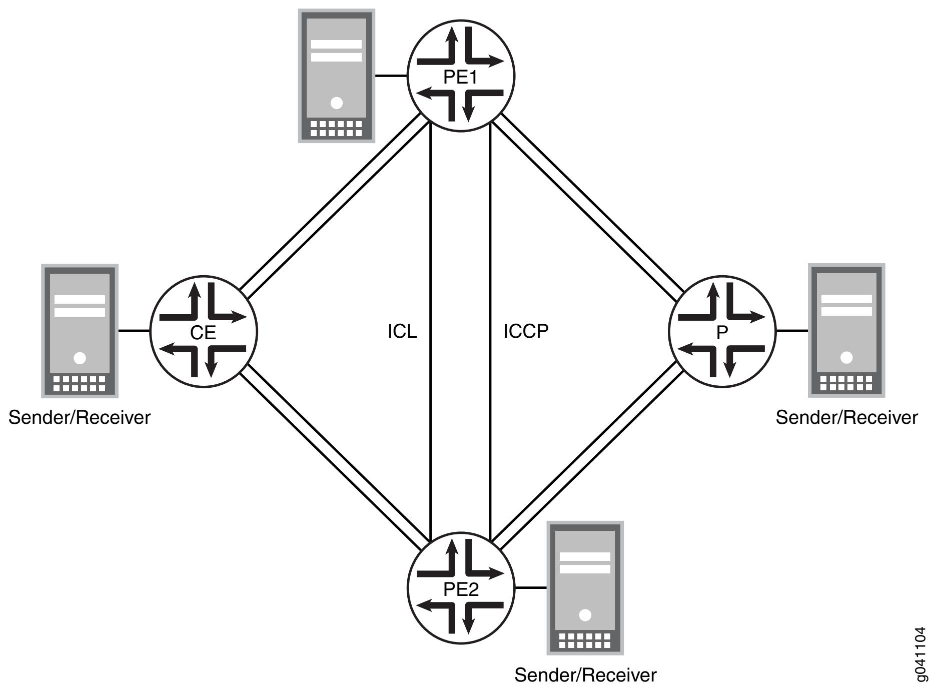

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/1/4.0 bfd-liveness-detection
  transmit-interval threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface ge-1/1/4.0 priority 200
set protocols pim interface ge-1/1/4.0 version 2
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface ge-1/1/4.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface ge-1/1/1.0 priority 600
set protocols pim interface ge-1/1/1.0 version 2
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
```

 Copyright © 2013, Juniper Networks, Inc.

```
set protocols pim interface ge-1/1/1.0 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.2
```

Copyright © 2013, Juniper Networks, Inc.

```
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
   minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
   transmit-interval minimum-interval 350
set protocols ospf area 0.0.0.0 interface ge-1/0/4.0 bfd-liveness-detection
   transmit-interval threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface ge-1/0/4.0 priority 200
set protocols pim interface ge-1/0/4.0 version 2
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection minimum-receive-interval
   700
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection transmit-interval
   minimum-interval 350
set protocols pim interface ge-1/0/4.0 bfd-liveness-detection transmit-interval threshold
   500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@PE1# set aggregated-devices ethernet device-count 5
   ```

2. Specify the members to be included within the aggregated Ethernet bundles.

   ```
   [edit interfaces]
   user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
   user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
   ```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

   ```
   [edit interfaces]
   user@PE1# set ge-1/1/1 flexible-vlan-tagging
   user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
   user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
   user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
   user@PE1# set ge-1/1/4 flexible-vlan-tagging
   user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
   user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
   user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
   user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
   ```

 Copyright © 2013, Juniper Networks, Inc.

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 118).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 118) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs on PE1 and PE2:

1. Create a routed VLAN interface (RVI) for each MC-LAG, assign a virtual IP address that is shared between each router in the VRRP groups, and assign an individual IP address for each router in the VRRP groups.

PE1

```
[edit interfaces]
user@PE1# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1
  virtual-address 10.1.1.1
```

 Copyright © 2013, Juniper Networks, Inc.

user@PE1# **set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 virtual-address 10.1.1.2**

PE2

[edit interfaces]
user@PE2# **set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 virtual-address 10.1.1.1**
user@PE2# **set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 virtual-address 10.1.1.2**

2. Assign the priority for each router in the VRRP groups:

NOTE: The router configured with the highest priority is the master.

PE1

[edit interfaces]
user@PE1# **set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority 200**
user@PE1 **set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority 200**

PE2

[edit interfaces]
user@sPE2# **set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority 150**
user@PE2# **set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority 150**

3. Enable the router to accept all packets destined for the virtual IP address if it is the master in a VRRP group:

PE1

[edit interfaces]
user@PE1# **set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 accept-data**
user@PE1# **set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 accept-data**

PE2

[edit interfaces]
user@PE2# **set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 accept-data**
user@PE2# **set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 accept-data**

Step-by-Step Procedure

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area on PE1 and PE2.

   [edit protocols ospf]
   user@host# **set area 0.0.0.0**

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area on PE1 and PE2.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface ge-1/1/1.0
user@host# set interface ge-1/4/1.0
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces on PE1 and PE2.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
  700
user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
  threshold 500
user@host# set interface ge-1/4/1.0 bfd-liveness-detection
  minimum-receive-interval 700
user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
  minimum-interval 350
user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
  threshold 500
```

Step-by-Step Procedure

To configure PIM as the multicast protocol:

1. Configure a static rendezvous point (RP) address on PE1 and PE2.

```
[edit protocols pim]
user@host# set rp static address 1.0.0.3
```

2. Configure the address ranges of the multicast groups for which PE1 and PE2 can be a rendezvous point (RP).

```
[edit protocols pim rp static address 1.0.0.3]
user@host# set group-ranges 239.0.0.0/8
```

3. Enable PIM on the VLAN interfaces for the MC-LAGs on PE1 and PE2.

```
[edit protocols pim]
user@host# set interface ge-1/1/1.0 version 2
user@host# set interface ge-1/4/1.0 version 2
```

4. Configure each PIM interface's priority for being selected as the designated router (DR).

   An interface with a higher priority value has a higher probability of being selected as the DR.

   PE1

```
[edit protocols pim]
user@host# set interface ge-1/1/1.0 priority 200
user@host# set interface ge-1/4/1.0 priority 600
```

   PE2

```
[edit protocols pim]
user@host# set interface ge-1/1/1.0 priority 100
user@host# set interface ge-1/4/1.0 priority 500
```

 Copyright © 2013, Juniper Networks, Inc.

5. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the PIM interfaces on PE1 and PE2.

   ```
   [edit protocols pim]
   user@host# set interface ge-1/1/1.0 bfd-liveness-detection minimum-receive-interval
     700
   user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
     minimum-interval 350
   user@host# set interface ge-1/1/1.0 bfd-liveness-detection transmit-interval
     threshold 500
   user@host# set interface ge-1/4/1.0 bfd-liveness-detection
     minimum-receive-interval 700
   user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
     minimum-interval 350
   user@host# set interface ge-1/4/1.0 bfd-liveness-detection transmit-interval
     threshold 500
   ```

Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on PE1 and PE2.

   ```
   [edit]
   user@host# set protocols rstp interface all mode point-to-point
   ```
2. Disable RSTP on the ICL-PL interfaces on PE1 and PE2:

   ```
   [edit]
   user@host# set protocols rstp interface ae0.0 disable
   ```
3. Configure the MC-LAG interfaces as edge ports on PE1 and PE2.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

   ```
   [edit]
   user@host# set protocols rstp interface ae1.0 edge
   ```
4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on PE1 and PE2.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

   ```
   [edit]
   user@host# set protocols rstp bpdu-block-on-edge
   ```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
}
```

```
user@PE1# show chassis
aggregated-devices {
  ethernet {
    device-count 5;
  }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/6 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-1/0/2 {
  unit 0 {
    family inet {
      address 100.100.100.1/30;
    }
  }
}
ge-1/1/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
ge-1/1/4 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
```

 Copyright © 2013, Juniper Networks, Inc.

```
            vlan-id-range 100-110;
        }
    }
    ae0 {
        flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        aggregated-ether-options {
            lacp {
                active;
                system-priority 100;
                system-id 00:00:00:00:00:05;
                admin-key 1;
            }
            mc-ae {
                mc-ae-id 5;
                redundancy-group 10;
                chassis-id 1;
                mode active-active;
                status-control active;
            }
        }
        unit 0 {
            encapsulation vlan-bridge;
            vlan-id-range 100-110;
            multi-chassis-protection 100.100.100.2 {
                interface ge-1/1/4.0;
            }
        }
    }
    ae1 {
        flexible-vlan-tagging;
        encapsulation flexible-ethernet-services;
        aggregated-ether-options {
            lacp {
                active;
                system-priority 100;
                system-id 00:00:00:00:00:05;
                admin-key 1;
            }
            mc-ae {
                mc-ae-id 10;
                redundancy-group 10;
                chassis-id 1;
                mode active-active;
                status-control active;
            }
        }
        unit 0 {
            encapsulation vlan-bridge;
            vlan-id-range 100-110;
            multi-chassis-protection 100.100.100.2 {
                interface ge-1/1/4.0;
            }
        }
    }
```

```
user@PE1#show vrrp
vlan {
   unit 100 {
      family inet {
         address 10.1.1.11/24 {
            vrrp-group 1 {
               virtual-address 10.1.1.1;
               priority 200;
               accept-data;
            }
         }
      }
   }
   unit 200 {
      family inet {
         address 10.1.1.21/24 {
            vrrp-group 2 {
               virtual-address 10.1.1.2;
               priority 200;
               accept-data;
            }
         }
      }
   }
}

user@PE1# show protocols
iccp {
   local-ip-addr 100.100.100.1;
   peer 100.100.100.2 {
      redundancy-group-id-list 10;
      liveness-detection {
         minimum-interval 1000;
      }
      rstp {
         interface ae0.0 {
            disable;
         }
         interface ae1.0 {
            edge;
         }
         interface all {
            mode point-to-point;
         }
         bpdu-block-on-edge;
      }
   }
}
rstp {
   interface ae0.0 {
      disable;
   }
   interface ae1.0 {
      edge;
   }
   interface all {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            mode point-to-point;
        }
        bpdu-block-on-edge;
    }
    ospf {
        area 0.0.0.0 {
            interface ge-1/1/1.0 {
                bfd-liveness-detection {
                    minimum-receive-interval 700;
                    transmit-interval {
                        minimum-interval 350;
                        threshold 500;
                    }
                }
            }
            interface ge-1/4/1.0 {
                bfd-liveness-detection {
                    minimum-receive-interval 700;
                    transmit-interval {
                        minimum-interval 350;
                        threshold 500;
                    }
                }
            }
        }
    }
    pim {
        rp {
            static {
                address 1.0.0.3 {
                    group-ranges {
                        239.0.0.0/8;
                    }
                }
            }
        }
        interface ge-1/1/1.0 {
            priority 200;
            version 2;
            bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
            }
        }
        interface ge-1/4/1.0 {
            priority 600;
            version 2;
            bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
                minimum-receive-interval 700;
                transmit-interval {
                    minimum-interval 350;
                    threshold 500;
                }
```

```
        }
      }
    }
    user@PE1# show switch-options
    service-id 10;
```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@CE# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
```

 Copyright © 2013, Juniper Networks, Inc.

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-2/1/6]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}
```

```
user@CE# show chassis
```

Copyright © 2013, Juniper Networks, Inc.

```
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/0/3 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/1/6 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
      }
   }
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
```


Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@P# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@P# set ge-1/0/5 gigether-options 802.3ad ae1
   user@P# set ge-1/0/11 gigether-options 802.3ad ae1
   ```

3. Configure an interface that connects to senders or receivers.

   ```
   [edit interfaces ge-1/1/3]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-500
   ```

4. Configure parameters on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1]
   user@P# set flexible-vlan-tagging
   user@P# set encapsulation flexible-ethernet-services
   user@P# set unit 0 encapsulation vlan-bridge
   user@P# set unit 0 vlan-id-range 100-110
   ```

5. Configure LACP on the aggregated Ethernet bundle.

   ```
   [edit interfaces ae1 aggregated-ether-options]
   user@P# set lacp active
   user@P# set lacp system-priority 100
   ```

6. Configure a domain that includes the set of logical ports.

   ```
   [edit bridge-domains bd0]
   user@P# set vlan-id all
   user@P# set domain-type bridge
   user@P# set interface ge-1/1/3.0
   ```

Copyright © 2013, Juniper Networks, Inc.

```
user@P# set interface ae1.0
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-1/1/3.0;
  interface ae1.0;
}

user@P# show chassis
aggregated-devices {
  ethernet {
    device-count 2;
  }
}

user@P# show interfaces
ge-1/0/5 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/0/11 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-1/1/3 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-500;
  }
}
ae1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  aggregated-ether-options {
    lacp {
      active;
      system-priority 100;
    }
  }
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 100-110;
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show pim interfaces**
- **show vrrp**
- ***show l2-learning instance* extensive**

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 98
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 110
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 115
- Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers on page 149

Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router

Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the routers participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 unicast MC-LAG is included in this example.

Requirements

This example uses the following hardware and software components:

- Four Juniper Networks MX Series routers.
- Junos OS Release 11.2 or later running on all four routers.

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a router.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a router.
- Configure Virtual Router Redundancy Protocol (VRRP) on a router.

Overview

In this example, you configure an MC-LAG across two routers by including interfaces from both routers in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.

NOTE: Layer 3 connectivity is required for ICCP.

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each router in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group

 Copyright © 2013, Juniper Networks, Inc.

Consider a sample topology in which a customer edge router, CE, is connected to two provider edge (PE) routers, PE1 and PE2, respectively. The two PE devices each have a LAG connected to the CE device. The configured mode is active -active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time. PE1 and PE2 are connected to a single service provider router, P.

In this example, the CE router is not aware that its aggregated Ethernet links are connected to two separate PE devices. The two PE devices each have a LAG connected to the CE device. The configured mode is active-active, meaning that both PE routers' LAG ports are active and carrying traffic at the same time.

In Figure 12 on page 117, from the perspective of Router CE, all four ports belonging to a LAG are connected to a single service provider device. Because the configured mode is active-active, all four ports are active, and the CE device load-balances the traffic to the peering PE devices. On the PE routers, a regular LAG is configured facindg the CE device.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG routers. Each of the routers has one or more physical links connected to a single client device. The routers coordinate with each other to ensure that data traffic is forwarded properly.

ICCP messages are sent between the two PE devices. In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, and ICCP for the peers hosting the MC-LAG.

Topology Diagram

Figure 12 on page 117 shows the topology used in this example.

Figure 14: MC-LAG Active-Active on MX Series Routers

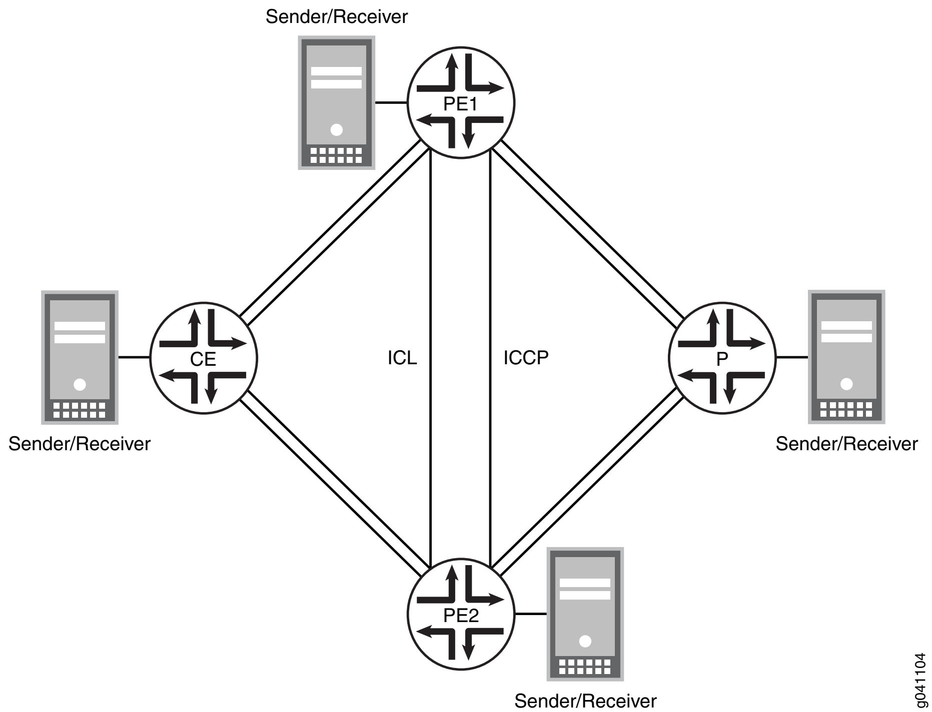

Configuring the PE Routers

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/1 gigether-options 802.3ad ae1
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.1/30
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/1/1 flexible-vlan-tagging
set interfaces ge-1/1/1 encapsulation flexible-ethernet-services
set interfaces ge-1/1/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/1 unit 0 vlan-id-range 100-110
set interfaces ge-1/1/4 flexible-vlan-tagging
set interfaces ge-1/1/4 encapsulation flexible-ethernet-services
set interfaces ge-1/1/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/4 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
```


Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/1/1.0
set bridge-domains bd0 interface ge-1/1/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.1
set protocols iccp peer 100.100.100.2 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.2 liveness-detection minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/0/2 unit 0 family inet address 100.100.100.2/30
set interfaces ge-1/0/3 flexible-vlan-tagging
set interfaces ge-1/0/3 encapsulation flexible-ethernet-services
set interfaces ge-1/0/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/3 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/4 flexible-vlan-tagging
set interfaces ge-1/0/4 encapsulation flexible-ethernet-services
set interfaces ge-1/0/4 unit 0 encapsulation vlan-bridge
set interfaces ge-1/0/4 unit 0 vlan-id-range 100-110
set interfaces ge-1/0/5 gigether-options 802.3ad ae0
set interfaces ge-1/1/0 gigether-options 802.3ad ae1
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae0 aggregated-ether-options lacp admin-key 1
set interfaces ae0 aggregated-ether-options mc-ae mc-ae-id 5
set interfaces ae0 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae0 aggregated-ether-options mc-ae chassis-id 0
```

```
set interfaces ae0 aggregated-ether-options mc-ae mode active-active
set interfaces ae0 aggregated-ether-options mc-ae status-control active
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-110
set interfaces ae0 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:05
set interfaces ae1 aggregated-ether-options lacp admin-key 1
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 10
set interfaces ae1 aggregated-ether-options mc-ae redundancy-group 10
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set interfaces ae1 unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 service-id 20
set bridge-domains bd0 interface ae1.0
set bridge-domains bd0 interface ge-1/0/3.0
set bridge-domains bd0 interface ge-1/0/4.0
set bridge-domains bd0 interface ae0.0
set protocols iccp local-ip-addr 100.100.100.2
set protocols iccp peer 100.100.100.1 redundancy-group-id-list 10
set protocols iccp peer 100.100.100.1 liveness-detection minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set switch-options service-id 10
```

Router PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@PE1# set aggregated-devices ethernet device-count 5
   ```

2. Specify the members to be included within the aggregated Ethernet bundles.

   ```
   [edit interfaces]
   user@PE1# set ge-1/0/1 gigether-options 802.3ad ae1
   user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
   ```

 Copyright © 2013, Juniper Networks, Inc.

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE1# set ge-1/1/1 flexible-vlan-tagging
user@PE1# set ge-1/1/1 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/1 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/1 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/1/4 flexible-vlan-tagging
user@PE1# set ge-1/1/4 encapsulation flexible-ethernet-services
user@PE1# set ge-1/1/4 unit 0 encapsulation vlan-bridge
user@PE1# set ge-1/1/4 unit 0 vlan-id-range 100-110
user@PE1# set ge-1/0/2 unit 0 family inet address 100.100.100.1/30
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0

[edit interfaces ae1]
user@PE1# set flexible-vlan-tagging
user@PE1# set encapsulation flexible-ethernet-services
user@PE1# set unit 0 encapsulation vlan-bridge
user@PE1# set unit 0 vlan-id-range 100-110
user@PE1# set unit 0 multi-chassis-protection 100.100.100.2 interface ge-1/1/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set lacp active
user@PE1# set lacp system-priority 100
user@PE1# set lacp system-id 00:00:00:00:00:05
user@PE1# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 5
user@PE1# set mc-ae redundancy-group 10
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active

[edit interfaces ae1 aggregated-ether-options]
user@PE1# set mc-ae mc-ae-id 10
user@PE1# set mc-ae redundancy-group 10
```

```
user@PE1# set mc-ae chassis-id 1
user@PE1# set mc-ae mode active-active
user@PE1# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**. (To refer to the configuration on Router PE2, see "Router PE2" on page 118).

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE1 uses **chassid-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE2 (as shown in "Router PE2" on page 118) uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE1# set domain-type bridge
user@PE1# set vlan-id all
user@PE1# set service-id 20
user@PE1# set interface ae0.0
user@PE1# set interface ae1.0
user@PE1# set interface ge-1/0/3.0
user@PE1# set interface ge-1/1/1.0
user@PE1# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE1# set local-ip-addr 100.100.100.1
user@PE1# set peer 100.100.100.2 redundancy-group-id-list 10
user@PE1# set peer 100.100.100.2 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE1# set service-id 10
```


Copyright © 2013, Juniper Networks, Inc.

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

Step-by-Step Procedure

To enable VRRP on the MC-LAGs on PE1 and PE2:

1. Enable VRRP on the MC-LAG on PE1 and PE2:

   - Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each router in the VRRP group, and assign an individual IP address for each router in the VRRP group:

     PE1:

     ```
     [edit interfaces]
     user@PE1# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 virtual-address 100.1.1.1
     ```

     PE2:

     ```
     [edit interfaces]
     user@PE2# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 virtual-address 100.1.1.1
     ```

   - Assign the priority for each router in the VRRP group:

     NOTE: The router configured with the highest priority is the master.

     PE1:

     ```
     [edit interfaces]
     user@PE1# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 priority 200
     ```

     PE2:

     ```
     [edit interfaces]
     user@PE2# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
     ```

   - Enable the router to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

     PE1:

     ```
     [edit interfaces]
     user@PE1# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 accept-data
     ```

     PE2:

     ```
     [edit interfaces]
     user@PE2# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
     ```

Step-by-Step Procedure

To enabled RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

   ```
   [edit]
   user@host# set protocols rstp interface all mode point-to-point
   ```

2. Disable RSTP on the ICL-PL interfaces on PE1 and PE2:

 Copyright © 2013, Juniper Networks, Inc.

```
[edit]
user@host# set protocols rstp interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@host# set protocols rstp interface ae1.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

   NOTE: The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@host# set protocols rstp bpdu-block-on-edge
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, **show interfaces**, **show protocols**, and **show switch-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  service-id 20;
  interface ae1.0;
  interface ge-1/1/1.0;
  interface ge-1/1/4.0;
  interface ae0.0;
}

user@PE1# show vrrp
vlan {
unit 100 {
  family inet {
    address 100.1.1.11/24 {
      vrrp-group 1 {
        virtual-address 100.1.1.1;
        priority 200;
        accept-data;
      }
    }
  }
}

user@PE1# show chassis
```

```
aggregated-devices {
   ethernet {
      device-count 5;
   }
}
```

```
user@PE1# show interfaces
ge-1/0/1 {
   gigether-options {
      802.3ad ae1;
   }
}
ge-1/0/6 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-1/0/2 {
   unit 0 {
      family inet {
         address 100.100.100.1/30;
      }
   }
}
ge-1/1/1 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ge-1/1/4 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
         system-id 00:00:00:00:00:05;
         admin-key 1;
      }
      mc-ae {
         mc-ae-id 5;
         redundancy-group 10;
         chassis-id 1;
         mode active-active;
         status-control active;
```

 Copyright © 2013, Juniper Networks, Inc.

```
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
            system-id 00:00:00:00:00:05;
            admin-key 1;
        }
        mc-ae {
            mc-ae-id 10;
            redundancy-group 10;
            chassis-id 1;
            mode active-active;
            status-control active;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
        multi-chassis-protection 100.100.100.2 {
            interface ge-1/1/4.0;
        }
    }
}
```

```
user@PE1# show protocols
iccp {
    local-ip-addr 100.100.100.1;
    peer 100.100.100.2 {
        redundancy-group-id-list 10;
        liveness-detection {
            minimum-interval 1000;
        }
        rstp {
            interface ae0.0 {
                disable;
            }
            interface ae1.0 {
                edge;
            }
            interface all {
                mode point-to-point;
            }
            bpdu-block-on-edge;
```

```
        }
      }
    }
user@PE1# show switch-options
service-id 10;
```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for Router PE2, using the appropriate interface names and addresses.

Configuring the CE Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router CE

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-2/0/2 gigether-options 802.3ad ae0
set interfaces ge-2/0/3 gigether-options 802.3ad ae0
set interfaces ge-2/1/6 flexible-vlan-tagging
set interfaces ge-2/1/6 encapsulation flexible-ethernet-services
set interfaces ge-2/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-2/1/6 unit 0 vlan-id-range 100-110
set interfaces ae0 flexible-vlan-tagging
set interfaces ae0 encapsulation flexible-ethernet-services
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp system-priority 100
set interfaces ae0 unit 0 encapsulation vlan-bridge
set interfaces ae0 unit 0 vlan-id-range 100-500
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 interface ge-2/1/6.0
set bridge-domains bd0 interface ae0.0
```

Router CE

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@CE# set aggregated-devices ethernet device-count 2
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces]
   user@CE# set ge-2/0/2 gigether-options 802.3ad ae0
   user@CE# set ge-2/0/3 gigether-options 802.3ad ae0
   ```


Copyright © 2013, Juniper Networks, Inc.

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-2/1/6]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-110
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae0]
user@CE# set flexible-vlan-tagging
user@CE# set encapsulation flexible-ethernet-services
user@CE# set unit 0 encapsulation vlan-bridge
user@CE# set unit 0 vlan-id-range 100-500
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae0 aggregated-ether-options]
user@CE# set lacp active
user@CE# set lacp system-priority 100
```

   The **active** statement initiates transmission of LACP packets.

   For the **system-priority** statement, a smaller value indicates a higher priority. The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby mode for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored.

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@CE# set domain-type bridge
user@CE# set vlan-id all
user@CE# set interface ge-2/1/6.0
user@CE# set interface ae0.0
```

   The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE# show bridge-domains
bd0 {
  domain-type bridge;
  vlan-id all;
  interface ge-2/1/6.0;
  interface ae0.0;
}
```

```
user@CE# show chassis
```

```
aggregated-devices {
   ethernet {
      device-count 2;
   }
}
```

```
user@CE# show interfaces
ge-2/0/2 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/0/3 {
   gigether-options {
      802.3ad ae0;
   }
}
ge-2/1/6 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-110;
   }
}
ae0 {
   flexible-vlan-tagging;
   encapsulation flexible-ethernet-services;
   aggregated-ether-options {
      lacp {
         active;
         system-priority 100;
      }
   }
   unit 0 {
      encapsulation vlan-bridge;
      vlan-id-range 100-500;
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the Provider Router

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router P

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ge-1/0/5 gigether-options 802.3ad ae1
set interfaces ge-1/0/11 gigether-options 802.3ad ae1
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id-range 100-500
```


Copyright © 2013, Juniper Networks, Inc.

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-priority 100
set interfaces ae1 unit 0 encapsulation vlan-bridge
set interfaces ae1 unit 0 vlan-id-range 100-110
set bridge-domains bd0 vlan-id all
set bridge-domains bd0 domain-type bridge
set bridge-domains bd0 interface ge-1/1/3.0
set bridge-domains bd0 interface ae1.0
```

Router P

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router P:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@P# set aggregated-devices ethernet device-count 2
```

2. Specify the members to be included within the aggregated Ethernet bundle.

```
[edit interfaces]
user@P# set ge-1/0/5 gigether-options 802.3ad ae1
user@P# set ge-1/0/11 gigether-options 802.3ad ae1
```

3. Configure an interface that connects to senders or receivers.

```
[edit interfaces ge-1/1/3]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-500
```

4. Configure parameters on the aggregated Ethernet bundle.

```
[edit interfaces ae1]
user@P# set flexible-vlan-tagging
user@P# set encapsulation flexible-ethernet-services
user@P# set unit 0 encapsulation vlan-bridge
user@P# set unit 0 vlan-id-range 100-110
```

5. Configure LACP on the aggregated Ethernet bundle.

```
[edit interfaces ae1 aggregated-ether-options]
user@P# set lacp active
user@P# set lacp system-priority 100
```

6. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@P# set vlan-id all
user@P# set domain-type bridge
user@P# set interface ge-1/1/3.0
```

Copyright © 2013, Juniper Networks, Inc.

```
user@P# set interface ae1.0
```

Results

From configuration mode, confirm your configuration by entering the **show bridge-domains**, **show chassis**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P# show bridge-domains
bd0 {
    domain-type bridge;
    vlan-id all;
    interface ge-1/1/3.0;
    interface ae1.0;
}

user@P# show chassis
aggregated-devices {
    ethernet {
        device-count 2;
    }
}

user@P# show interfaces
ge-1/0/5 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/0/11 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-1/1/3 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-500;
    }
}
ae1 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-priority 100;
        }
    }
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id-range 100-110;
    }
}
```


Copyright © 2013, Juniper Networks, Inc.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly by running the following commands:

- **show iccp**
- **show interfaces ae0**
- **show interfaces ae1**
- **show interfaces mc-ae**
- **show vrrp**
- ***show l2-learning instance* extensive**

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the MC-LAG is **down**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

Related Documentation

- Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview on page 98
- Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers on page 110
- Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers on page 115
- Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using Virtual Router Redundancy Protocol (VRRP) on MX Series Routers on page 129

IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview

IGMP Snooping in MC-LAG Active-Active on MX Series Routers Functionality

MX Series routers support multichassis link aggregation group (MC-LAG) active-active and IGMP snooping in active-standby mode. MC-LAG allows one device to form a logical LAG interface with two or more network devices. MC-LAG provides additional benefits including node level redundancy, multi-homing, and loop-free layer-2 network without running STP. The following features are supported:

- State synchronization between peers for IGMP snooping in a bridge domain with only Layer 2 interfaces
- Qualified learning
- Router facing multichassis links

MX Series routers support the following enhancements to active-active bridging and virtual router redundancy protocol (VRRP) over integrated routing and bridging (IRB):

- MC-LAG support for IGMP snooping in a pure Layer 2 switch
- MC-LAG support for IGMP snooping in bridge domains doing qualified learning
- Support for MC-Links being router facing interfaces

The following functions are not supported:

- MC-LAG for VPLS instances
- MC-Links trunk ports
- Proxy mode for active-active
- Adding interchassis links to outgoing interfaces on an as needed basis. Interchassis links can be added to the outgoing interface list as router facing interfaces.

 Copyright © 2013, Juniper Networks, Inc.

Typically Supported Network Topology for IGMP Snooping with MC-LAG Active-Active Bridging

Figure 15 on page 169 depicts a typical network topology over which IGMP snooping with MC-LAG active-active is supported.

Figure 15: Typical Network Over Which Active-Active Is Supported

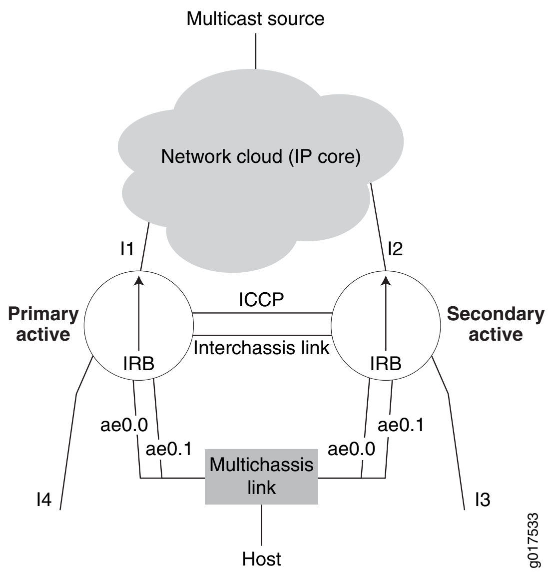

Interfaces I3 and I4 are single-homed interfaces. The multichassis links (MC-Link) ae0.0 and ae0.1 belong to the same bridge domain in both the chassis. Interfaces I3, ae0.0 and ae0.1 are in the same bridge domain in S-A. Interfaces I4, ae0.0 and ae0.1 are in the same bridge domain in the primary active (P-A) router. Interfaces I3, I4, ae0.0 and ae0.1 are in the same learning domain as is the interchassis link (ICL) connecting the two chassis.

The primary active router is the chassis in which the integrated routing and bridging has become PIM-DR. The secondary active router is the chassis in which integrated routing and bridging is not PIM DR. Router P-A is the chassis responsible for pulling traffic from the IP core. Hence, PIM-DR election is used to avoid duplication of data traffic.

Learning domains are described in "Qualified Learning" on page 171.

For the IGMP speakers (hosts and routers) in the learning domain, P-A and S-A together should appear as one device with interfaces I4, I3, ae0.0 and ae0.1.

No duplicate control packets should be sent on multichassis links, meaning the control packet should be sent through only one link.

Control Plane State Updates Triggered by Packets Received on Remote Chassis

The membership state in Layer 3 multicast routing is updated as a result of reports learned on remote legs of multichassis links and s-links attached to the remote chassis.

The membership state and routing entry in snooping is updated when reports are received on the remote legs of a multichassis link.

When reports are received on S-links attached to the remote chassis the membership state or routing entry in snooping is not updated.

The list of <s,g>s for which the state is maintained is the same in both the chassis under snooping as long as the outgoing interface lists involve only multichassis links.

Data Forwarding

This discussion assumes integrated routing and bridging on P-A is the PIM-DR. It pulls the traffic from sources in the core. Traffic might also come on Layer 2 interfaces in the bridge domain. For hosts directly connected to the P-A chassis, there is no change in the way data is delivered.

For delivering traffic to hosts connected to S-A (which is the non-DR) on the single-homed link like l3, we rely on interchassis link. The traffic that hits P-A is sent over ICL to S-A to be delivered to the links that have reported interests in s,g and the links that are router facing.

When ae0 leg in P-A goes down, the hosts connected to the multichassis link will receive traffic via ICL. In S-A, traffic received on ICL is sent to multichassis links in the outgoing interface list for which the ae counterpart in P-A is down.

 Copyright © 2013, Juniper Networks, Inc.

Pure Layer 2 Topology Without Integrated Routing and Bridging

Figure 16 on page 171 illustrates the chassis connecting to the PIM-DR is the primary active router and the other is the secondary active.

Figure 16: Layer 2 Configuration Without Integrated Routing and Bridging

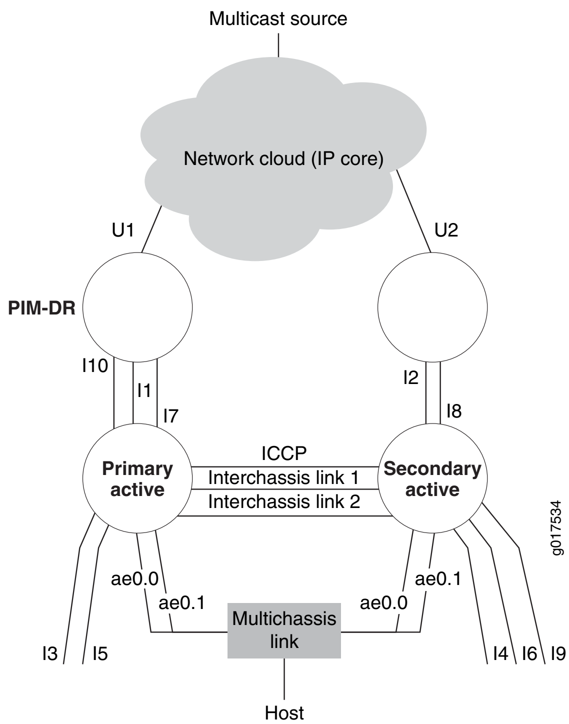

Qualified Learning

In this application, interfaces I1, I2, I3, I4, I5, I6, I7, I8, I9 and I10 are single-homed interfaces. The multichassis links ae0.0 and ae0.1 belong to the same bridge domain in both the chassis. Interfaces I10, I1,I7,I3,I5, ae0.0 and ae0.1 are in same bridge domain, bd1 in P-A. Interfaces I9, I2,I8,I4,I6, ae0.0 and ae0.1 are in same bridge domain, bd1 in S-A.

This discussion assumes the following configuration:

- In Primary Active and S-A, qualified learning is ON in bd1.
- Interfaces I1,I2,I3, ae0.0 and I4 belong to vlan1, learning domain ld1.
- Interfaces I7,I8,I5, ae0.1 and I6 belong to vlan2, learning domain ld2.
- Interfaces I9 and I10 belong to vlan3, learning domain ld3.

For the IGMP speakers (hosts and routers) in the same learning domain ld1, P-A and S-A linked should appear to be one switch.

For the IGMP speakers (hosts and routers) in the same learning domain ld2, P-A and S-A linked should appear to be one switch.

Copyright © 2013, Juniper Networks, Inc.

Since there are no multichassis links in learning domain ld3, for the IGMP speakers (hosts and routers) in learning domain ld3, P-A and S-A will not appear to be one switch.

This discussion assumes interchassis link ICL1 corresponds to learning domain ld1 and interchassis link ICL2 corresponds to learning domain ld2.

Control packet flow is supported, with the exception of passing information to IRB.

Data Forwarding with Qualified Learning

This discussion assumes one learning domain (LD), ld1, and further assumes interface I1 on router P-A is connected to the PIM-DR in the learning domain and pulls the traffic from sources in the core.

For delivering traffic to hosts connected to router S-A (which is the non-DR) on the single-homed link like I2, I4 (belonging to ld1), we rely on ICL1. The traffic that hits router P-A on interface I1 is sent over interchassis link ICL1 to router S-A to be delivered to the links that have reported interests in s,g or the links that are router facing in learning domain ld1.

When the interface ae0 leg in router P-A goes down, the hosts connected to the multichassis link receive traffic from interface I1 via the interchassis link ICL1. In router S-A, traffic received on interchassis link ICL1 is sent to multichassis links in the outgoing interface list for which the aggregated Ethernet counterpart in router P-A is down.

It is further assumed that interface I9 in router S-A belongs to the learning domain ld3 with interests in s,g, and that interface I10 in learning domain ld3 in router P-A receives traffic for s,g. Interface I9 does not receive data in this topology because there are no multichassis links (in a-a mode) and hence no interchassis link in learning domain ld3.

Static Groups on Single Homed Interfaces

For multichassis links, the static group configuration should exist on both legs and synchronization with the other chassis is not required.

Synchronization of the static groups on single homed interfaces between the chassis is not supported, however the addition of logical interfaces to the default outgoing interface list supports traffic delivery to the interface within a static configuration.

Router Facing Interfaces as Multichassis Links

IGMP queries could arrive on either leg of the multichassis links but in both peers, the multichassis link should be considered as router facing.

Reports should exit only once from the multichassis link, that is from only one leg.

The following MC-LAG support for IGMP snooping in IRB is provided:

- Non-proxy snooping
- Logical interfaces must be outgoing interfaces for all routes including the default route
- IGMP snooping in a pure Layer 2 switch

 Copyright © 2013, Juniper Networks, Inc.

- IGMP snooping in bridge domains doing qualified learning
- Router facing interface MC-Links

The following features are not supported:

- Proxy mode for active-active
- MC-LAG support for VPLS instances
- Trunk ports as multichassis links
- Adding logical interfaces to outgoing interfaces on need basis. However, logical interfaces are always added as a router facing interface to the outgoing interface list.

Related Documentation

- Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers on page 173
- *Example: Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers*
- *Example: Configuring IGMP Snooping*
- *igmp-snooping*
- *multicast-router-interface*
- *show l2-learning instance*
- *Ethernet Interfaces*

Configuring IGMP Snooping in MC-LAG Active-Active on MX Series Routers

You can use the bridge-domain statement's service-id id option to specify the multichassis aggregated Ethernet configuration.

- The **service-id** statement is mandatory for non-single VLAN type bridge domains (**none**, **all** or **vlan-id-tags:dual**).
- It is optional for bridge domains with a VID defined.
- If no service-id is defined in the latter case, it will be picked up from the RTT's **service-id** configuration.
- The bridge level service-id is required to link related bridge domains across peers, and should be configured with the same value.
- The service-id values share the name space across all bridging and routing instances, and across peers. Thus, duplicate values for service-ids are not permitted across these entities.
- A change of bridge **service-id** is considered catastrophic, and the bridge domain is reincarnated.

This procedure allows you to enable or disable the replication feature. This option applies to all instances.

Copyright © 2013, Juniper Networks, Inc.

To configure IGMP snooping in active-standby mode:

1. Use the **multichassis-lag-replicate-state** statement at the **multicast-snooping-options** hierarchy level in the master instance.

```
multicast-snooping-options {
  ...
  multichassis-lag-replicate-state; # REQUIRED
}
```

The interchassis link, **interface *icl-intf-name***, of the learning domain should be a router facing interface.

1. Use the **interface *icl-intf-name*** statement at the **protocols igmp-snooping** hierarchy level, as shown in the following example:

```
protocols {
  igmp-snooping {
    interface icl-intf-name {
      multicast-router-interface;
    }
  }
}
```

Related Documentation

- IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview on page 167
- *Example: Configuring IGMP Snooping*
- *igmp-snooping*
- *multicast-router-interface*
- *show l2-learning instance*
- *Ethernet Interfaces*

Configuring Aggregated Ethernet Link Protection

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

 Copyright © 2013, Juniper Networks, Inc.

NOTE: Link protection is not supported on MX80.

Configuring Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Configure the link protection mode.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **set link-protection**

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* primary**

2. Configure the backup logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* backup**

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

Copyright © 2013, Juniper Networks, Inc.

user@host> **request interface revert ae*x***

Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert ae*x*** configuration command.

user@host# **delete interfaces ae*x* aggregated-ether-options link-protection**

Configuring Shared Scheduling on Aggregated Ethernet Interfaces

You can configure shared scheduling on aggregated Ethernet interfaces in link protection mode on Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Ethernet Enhanced IQ2 (IQ2E) PICs on M320 routers.

To configure shared scheduling on aggregated Ethernet interfaces:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Configure the link protection mode.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **set link-protection**

3. Configure shared scheduling.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **top**
   [edit]
   user@host# **edit interfaces ae*x* shared-scheduler**

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 174

Configuring the Number of Aggregated Ethernet Interfaces on the Device

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them.

On M Series and T Series routers, you can configure a maximum number of 128 aggregated interfaces, whereas on MX Series routers you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces are numbered from **ae0** through **ae127** for M Series and T Series routers and the aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers.

1. Specify that you want to access the aggregated Ethernet configuration on the device.

   user@host# **edit chassis aggregated-devices ethernet**

2. Set the number of aggregated Ethernet interfaces.

   [edit chassis aggregated-devices ethernet]
   user@host# **set device-count *number***

Copyright © 2013, Juniper Networks, Inc.

You must also specify the constituent physical links by including the **802.3ad** statement at the **[edit interfaces *interface-name* fastether-options]** or **[edit interfaces *interface-name* gigether-options]** hierarchy level.

For information about E Series routers, see *Understanding Aggregated Ethernet Interfaces and LACP*.

Related Documentation

- For information about physical links, see Configuring an Aggregated Ethernet Interface on page 87
- For a sample configuration, see Example: Configuring Aggregated Ethernet Interfaces on page 544
- *Ethernet Interfaces*
- For information about configuring aggregated devices, see the *Junos OS Administration Library for Routing Devices*.

Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled.

For Multichassis Link Aggregation (MC-LAG), you must specify the **system-id** and **admin key**. MC-LAG peers use the same **system-id** while sending the LACP messages. The **system-id** can be configured on the MC-LAG network device and synchronized between peers for validation.

LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

The LACP mode can be active or passive. If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is turned off on aggregated Ethernet interfaces. If LACP is configured, it is in passive mode by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode.

Copyright © 2013, Juniper Networks, Inc.

To enable LACP active mode, include the **lacp** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **active** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
  active;
}
```

NOTE: The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

To restore the default behavior, include the **lacp** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **passive** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
  passive;
}
```

Starting with Junos OS Release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.

To override the IEEE 802.3ad standard and facilitate subsecond failover, include the **fast-failover** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you use the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

For more information, see the following sections:

- Configuring the LACP Interval on page 179
- Configuring LACP Link Protection on page 179

 Copyright © 2013, Juniper Networks, Inc.

Configuring the LACP Interval

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options lacp]
periodic interval;
```

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

NOTE: Source address filtering does not work when LACP is enabled. This behavior is not applicable to T Series routers and PTX Series Packet Transport Routers. For more information about source address filtering, see "Enabling Ethernet MAC Address Filtering" on page 42.

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the *Routing Policy Feature Guide for Routing Devices*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see "Configuring Untagged Aggregated Ethernet Interfaces" on page 184.

For M Series Multiservice Edge Routers with enhanced Flexible PIC Concentrators (FPCs) and T Series routers, LACP over VLAN-tagged aggregated Ethernet interfaces is supported. For 8-port, 12-port, and 48-port Fast Ethernet PICs, LACP over VLAN-tagged interfaces is not supported.

LACP **Fast Periodic**, which is achieved by configuring fast (every second) intervals for periodic transmission of LACP packets, is supported with graceful Routing Engine switchover (GRES) on MX Series routers only.

Configuring LACP Link Protection

NOTE: When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the **link-protection** and **system-priority** statements. Configuring values at this level results in

Copyright © 2013, Juniper Networks, Inc.

only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

NOTE: **LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.**

Enabling LACP Link Protection

To enable LACP link protection for an aggregated Ethernet interface, use the **link-protection** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable;
  revertive;
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, once a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch and the current link remains active.

If LACP link protection is configured to be nonrevertive at the global (**[edit chassis]** hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.

CAUTION: **If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.**

We strongly recommend that you use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces of MX Series routers to any other vendor device. Otherwise, the vendor device (say a Layer 2 switch, or a router) will not be able to manage the traffic coming from the two link aggregated Ethernet bundle. As a result, you might observe the vendor device sending back the traffic to the backup member link of the aggregated Ethernet interface.

Copyright © 2013, Juniper Networks, Inc.

Currently, MX-MPC2-3D, MX-MPC2-3D-Q, MX-MPC2-3D-EQ, MX-MPC1-3D, MX-MPC1-3D-Q, and MPC-3D-16XGE-SFPP do not drop traffic coming back to the backup link, whereas DPCE-R-Q-20GE-2XGE, DPCE-R-Q-20GE-SFP, DPCE-R-Q-40GE-SFP, DPCE-R-Q-4XGE-XFP, DPCE-X-Q-40GE-SFP, and DPCE-X-Q-4XGE-XFP drop traffic coming to the backup link.

Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 to 65,535.

Configuring LACP System Identifier

To configure the LACP system identifier for aggregated Ethernet interfaces, use the **system-id** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-id system-id;
```

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Configuring LACP administrative Key

To configure an administrative key for LACP, include the **admin-key *number*** statement at the **[edit interfaces ae*x* aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces ae x aggregated-ether-options-lacp]
admin-key number;
```

NOTE: You must configure MC-LAG to configure the **admin-key** statement. For more information about MC-LAG, see “Configuring Multichassis Link Aggregation” on page 93.

Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the **port-priority** statement at the **[edit interfaces *interface-name* gigether-options 802.3ad ae*X* lacp]** or **[edit interfaces *interface-name* fastether-options 802.3ad ae*X* lacp]** hierarchy levels:

```
[edit interfaces interface-name gigether-options 802.3ad aeX lacp]
port-priority priority;
```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and is assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.

NOTE: **Port aggregation selection (discussed previously) is performed for the active link when LACP link protection is enabled. Without LACP link protection, port priority is not used in port aggregation selection.**

Tracing LACP Operations

To trace the operations of the LACP process, include the **traceoptions** statement at the **[edit protocols lacp]** hierarchy level:

```
[edit protocols lacp]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

You can specify the following flags in the **protocols lacp traceoptions** statement:

Copyright © 2013, Juniper Networks, Inc.

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

For general information about tracing, see the tracing and logging information in the *Junos OS Administration Library for Routing Devices*.

LACP Limitations

LACP can link together multiple different physical interfaces, but only features that are supported across all of the linked devices will be supported in the resulting link aggregation group (LAG) bundle. For example, different PICs can support a different number of forwarding classes. If you use link aggregation to link together the ports of a PIC that supports up to 16 forwarding classes with a PIC that supports up to 8 forwarding classes, the resulting LAG bundle will only support up to 8 forwarding classes. Similarly, linking together a PIC that supports WRED with a PIC that does not support it will result in a LAG bundle that does not support WRED.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

LACP with VLAN-Tagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
  }
}
```

Configure aggregated Ethernet LACP over an untagged interface:

LACP with Untagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}
```

Related Documentation

- lacp on page 589
- link-protection on page 595
- *traceoptions*
- *Ethernet Interfaces*

Configuring Untagged Aggregated Ethernet Interfaces

When you configure an untagged Aggregated Ethernet interface, the existing rules for untagged interfaces apply. These rules are as follows:

- You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.
- You cannot include the **vlan-id** statement in the configuration of the logical interface.

Table 8 on page 185 lists untagged aggregated Ethernet and LACP support by PIC and router.

Copyright © 2013, Juniper Networks, Inc.

Table 8: Untagged Aggregated Ethernet and LACP Support by PIC and Platform

PIC Type	M Series	LACP	T Series	LACP
4-port Fast Ethernet PIC Type 1	Yes	Yes	Yes	Yes
1-port Gigabit Ethernet PIC Type 1	Yes	Yes	Yes	Yes
2-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
4-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
1-port 10-Gigabit Ethernet M160	Yes	Yes	NA	NA
10-port Gigabit Ethernet PIC Type 3	Yes (M120, M320)	Yes	Yes	Yes
1-port 10-Gigabit Ethernet PIC Type 3	N/A	NA	Yes	Yes
8-port Gigabit Ethernet PIC Type 3	Yes	Yes	Yes	Yes

The 8–port Fast Ethernet PIC does not support untagged aggregated Ethernet or LACP.

Syslog messages are logged if you try to configure an untagged aggregated Ethernet interface using an unsupported PIC type.

For more information about configuring LACP, see “Configuring Aggregated Ethernet LACP” on page 177.

Example: Configuring Untagged Aggregated Ethernet Interfaces

Configure an untagged aggregated Ethernet interface by omitting the **vlan-tagging** and **vlan-id** statements from the configuration:

```
[edit interfaces]
fe-5/0/1 {
   fastether-options {
      802.3ad ae0;
   }
}
ae0 {
   # vlan-tagging; OMIT FOR UNTAGGED AE CONFIGURATIONS
   unit 0 {
      # vlan-id 100; OMIT FOR UNTAGGED AE CONFIGURATIONS
      family inet {
         address 13.1.1.2/24 {
            vrrp-group 0 {
               virtual-address 13.1.1.4;
               priority 200;
            }
         }
      }
```

```
    }
  }
```

Related Documentation

- For more information about configuring LACP, see Configuring Aggregated Ethernet LACP on page 177.
- *Ethernet Interfaces*

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers, you need to configure the **mixed** option for the [**edit interfaces aex aggregated-ether-options link-speed**] statement.

To set the required link speed:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces** ***interface-name*** aggregated-ether-options

2. Configure the link speed.

   [edit interfaces *interface-name* aggregated-ether-options]

 Copyright © 2013, Juniper Networks, Inc.

user@host# **set link-speed** ***speed***

speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet interfaces on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.
- **50g**—Links are 50 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

Related Documentation

- **aggregated-ether-options on page 555**
- Configuring Mixed Aggregated Ethernet Links on page 91
- *Ethernet Interfaces*

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces *interface-name*** aggregated-ether-options

2. Configure the minimum number of links.

   [edit interfaces *interface-name* aggregated-ether-options]
   user@host# **set** minimum-links ***number***

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links *number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

Related Documentation

- aggregated-ether-options on page 555
- minimum-links on page 606

- *Ethernet Interfaces*

Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces

T Series and TX Matrix routers support multicast statistics collection on aggregated Ethernet interfaces in both ingress and egress directions. The multicast statistics functionality can be configured on a physical interface thus enabling multicast accounting for all the logical interfaces below the physical interface.

The multicast statistics information is displayed only when the interface is configured with the **multicast-statistics** statement, which is not enabled by default.

Multicast statistics collection requires at least one logical interface is configured with family inet or inet6; otherwise, the commit for **multicast-statistics** will fail.

 Copyright © 2013, Juniper Networks, Inc.

The multicast in/out statistics can be obtained via interfaces statistics query through CLI and via MIB objects through SNMP query.

To configure multicast statistics:

1. Include the **multicast-statistics** statement at the **[edit interfaces interface-name]** hierarchy level.

An example of a multicast statistics configuration for an aggregated Ethernet interface follows:

```
[edit interfaces]
ae0 {
  multicast-statistics;
}
```

To display multicast statistics, use the **show interfaces *interface-name* statistics detail** command.

Related Documentation

- *multicast-statistics*
- Configuring Multicast Statistics Collection on Ethernet Interfaces on page 51
- *Ethernet Interfaces*

Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection

On aggregated Ethernet interfaces, you can configure scheduler in non–link-protect mode on the following platforms:

- MX-Series
- M120 and M320 with IQ2 PIC
- T-series platforms (T620 and T320) with IQ2 PIC

The scheduler functions supported are:

- Per unit scheduler
- Hierarchical scheduler
- Shaping at the physical interface

To configure the hierarchical scheduler on aggregated Ethernet interfaces in the non link-protect mode, include the **hierarchical-scheduler** statement at the **[edit interfaces ae*X*]** hierarchy level:

```
[edit interfaces aeX hierarchical-scheduler]
```

Prior to Junos OS Release 9.6, the hierarchical scheduler mode on these models required the **aggregated-ether-options** statement **link-protection** option. If a **link-protection** option is not specified, the scheduler is configured in non-link-protect mode.

To specify the member link bandwidth derivation based on the equal division model **(scale)** or the replication model **(replicate)** on aggregated Ethernet interfaces, include the **member-link-scheduler (scale | replicate)** option at the **[edit class-of-service interfaces ae*X*]** hierarchy level. The default setting is **scale**.

[edit class-of-service interfaces ae*X* member-link-scheduler (scale | replicate)]

NOTE: In link-protect mode, only one link is active at a time and the other link acts as the backup link, whereas in a non link-protect mode, all the links of the aggregate bundle are active at the same time. There is no backup link. If a link goes down or a new link is added to the bundle, traffic redistribution occurs.

Related Documentation

- *Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links*
- *Ethernet Interfaces*
- For more information on the hierarchical scheduler (CoS), see the *Junos OS Class of Service Library for Routing Devices*.

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

This section describes configuration of symmetrical load balancing on an 802.3ad link aggregation group (LAG) on MX Series routers.

Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview

MX Series routers with Aggregated Ethernet PICs support symmetrical load balancing on an 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through deep packet inspection (DPI) devices over an LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Without symmetrical load balancing on an 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. By using this feature, a given flow of traffic (duplex) is ensured for the same devices in both directions.

Symmetrical load balancing on an 802.3ad LAG utilizes a mechanism of interchanging the source and destination addresses for a hash computation of fields, such as source address and destination address. The result of a hash computed on these fields is used to choose the link of the LAG. The hash-computation for the forward and reverse flow must be identical. This is achieved by swapping source fields with destination fields for the reverse flow. The swapped operation is referred to as *complement hash computation* or **symmetric-hash complement** and the regular (or unswapped) operation as

 Copyright © 2013, Juniper Networks, Inc.

symmetric-hash computation or **symmetric-hash**. The swappable fields are MAC address, IP address, and port.

Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

You can specify whether symmetric hash or complement hash is done for load-balancing traffic. To configure symmetric hash, use the **symmetric-hash** statement at the **[edit forwarding-options hash-key family inet]** hierarchy level. To configure symmetric hash complement, use the **symmetric-hash complement** statement and option at the **[edit forwarding-options hash-key family inet]** hierarchy level.

These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the **symmetric-hash** or **symmetric-hash complement** statement at the **[edit chassis hash-key family inet]** and **[edit chassis hash-key family multiservice]** hierarchy levels.

Consider the example in Figure 17 on page 191.

Figure 17: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

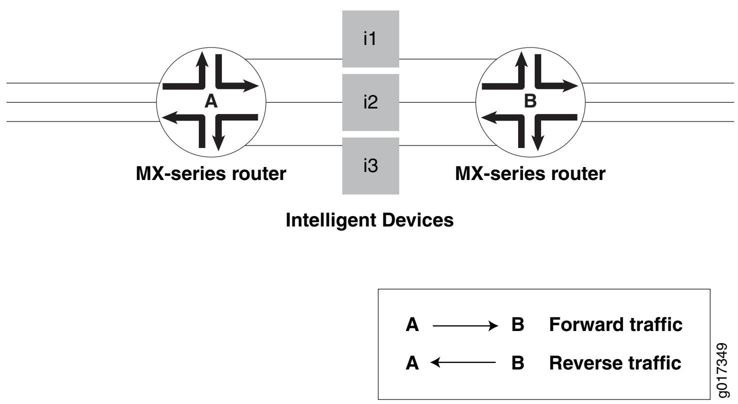

Router A is configured with symmetric hash and Router B is configured with symmetric hash complement. Thus, for a given flow *fx*, post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow *fx* is from Router B to Router A through the same i2 device as its hashing (done after swapping source and destination fields) and returns the same link index; since it is performed on the interchanged source and destination addresses.

However, the link chosen may or may not correspond to what was attached to the DPI. In other words, the hashing result should point to the same links that are connected, so that the traffic flows through the same DPI devices in both directions. To make sure this happens, you need to also configure the counterpart ports (ports that are connected to same DPI-iN) with the identical link index. This is done when configuring a child-link into the LAG bundle. This ensures that the link chosen for a given hash result is always the same on either router.

Note that any two links connected to each other should have the same link index and these link indices must be unique in a given bundle.

Copyright © 2013, Juniper Networks, Inc.

NOTE:

The following restrictions apply when configuring symmetric load balancing on an 802.3ad LAG on MX Series routers:

- The Packet Forwarding Engine (PFE) can be configured to hash the traffic in either symmetric or complement mode. A single PFE complex cannot work simultaneously in both operational modes and such a configuration can yield undesirable results.
- The per-PFE setting overrides the chassis-wide setting only for the family configured. For the other families, the PFE complex still inherits the chassis-wide setting (when configured) or the default setting.
- Any change in the hash key configuration requires a reboot of the FPC for the changes to take effect.
- This feature supports VPLS, INET, and bridged traffic only.
- This feature cannot work in tandem with the **per-flow-hash-seed load-balancing** option. It requires that all the PFE complexes configured in complementary fashion share the same seed. A change in the seed between two counterpart PFE complexes may yield undesired results.

For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library for Routing Devices*.

Example Configuration Statements

To configure 802.3ad LAG parameters at the bundle level:

```
[edit interfaces]
g(x)e-fpc/pic/port {
   gigether-options {
      802.3ad {
         bundle;
         link-index number;
      }
   }
}
```

where the **link-index *number*** ranges from 0 through 15.

You can check the link index configured above using the **show interfaces** command:

```
[edit forwarding-options hash-key]
family inet {
   layer-3;
   layer-4;
   symmetric-hash {
      [complement;]
   }
}
family multiservice {
   source-mac;
   destination-mac;
   payload {
```

 Copyright © 2013, Juniper Networks, Inc.

```
    ip {
      layer-3 {
        source-ip-only | destination-ip-only;
      }
      layer-4;
    }
    }
    symmetric-hash {
      [complement;]
    }
  }
```

For load-balancing Layer 2 traffic based on Layer 3 fields, you can configure 802.3ad LAG parameters at a per PIC level. These configuration options are available under the chassis hierarchy as follows:

```
[edit chassis]
fpc X {
  pic Y {
    .
    .
    .
    hash-key {
      family inet {
        layer-3;
        layer-4;
        symmetric-hash {
          [complement;]
        }
      }
      family multiservice {
        source-mac;
        destination-mac;
        payload {
          ip {
            layer-3 {
              source-ip-only | destination-ip-only;
            }
            layer-4;
          }
        }
        symmetric-hash {
          [complement;]
        }
      }
    }
    .
    .
    .
  }
}
```

Configuring Symmetrical Load Balancing on Trio-Based MPCs

With some configuration differences, symmetrical load-balancing over an 802.3ad link aggregation group is supported on MX Series routers with Trio-based MPCs.

Copyright © 2013, Juniper Networks, Inc.

To achieve symmetrical load-balancing on Trio-Based MPCs, the following needs to be done:

- Compute a Symmetrical Hash

  Both routers must compute the same hash value from the flow in the forward and reverse directions.On Trio-based platforms, the calculated hash value is independent of the direction of the flow, and hence is always symmetric in nature. For this reason, no specific configuration is needed to compute a symmetric hash value on Trio-based platforms.

  However, it should be noted that the fields used to configure the hash should have identical include and exclude settings on both ends of the LAG.

- Configure Link Indexes

  To allow both routers to choose the same link using the same hash value, the links within the LAG must be configured with the same link index on both routers. This can be achieved with the **link-index** statement.

- Enable Symmetric Load Balancing

  To configure symmetric load balancing on Trio-based MPCs, include the **symmetric** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy level. This statement is applicable to Trio-based platforms only.

  The **symmetric** statement can be used with any protocol family and enables symmetric load-balancing for all aggregated Ethernet bundles on the router. The statement needs to be enabled at both ends of the LAG. This statement is disabled by default.

- Achieve Symmetry for Bridged and Routed Traffic

  In some deployments, the LAG bundle on which symmetry is desired is traversed by Layer 2 bridged traffic in the upstream direction and by IPv4 routed traffic in the downstream direction. In such cases, the computed hash is different in each direction because the Ethernet MAC addresses are taken into account for bridged packets. To overcome this, you can exclude source and destination MAC addresses from the enhanced-hash-key computation.

  To exclude source and destination MAC addresses from the enhanced-hash-key computation, include the **no-mac-addresses** statement at the **[edit forwarding-options enhanced-hash-key family multiservice]** hierarchy level. This statement is disabled by default.

When symmetrical load balancing is enabled on Trio-based MPCs, keep in mind the following caveats:

- Traffic polarization is a phenomenon that occurs when using topologies that distribute traffic by using hashing of the same type. When routers are cascaded, traffic polarization can occur, and this can lead to unequal traffic distribution.

  Traffic polarization occurs when LAGs are configured on cascaded routers. For example, in Figure 18 on page 195, if a certain flow uses Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, the flow also uses Link 1 of the aggregated Ethernet bundle between Device R2 and Device R3.

 Copyright © 2013, Juniper Networks, Inc.

Figure 18: Traffic Polarization on Cascaded Routers When Symmetrical Load Balancing in Enabled on Trio-based MPCs

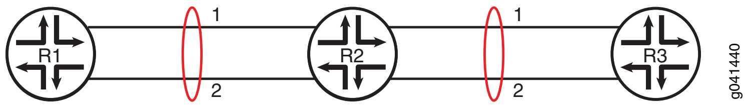

This is unlike having a random link selection algorithm, where a flow might use Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, and Link 2 of the aggregated Ethernet bundle between Device R2 and Device R3.

- Symmetric load balancing is not applicable to per-prefix load-balancing where the hash is computed based on the route prefix.
- Symmetric load balancing is not applicable to MPLS or VPLS traffic, because in these scenarios the labels are not the same in both directions.

Example Configurations

Example Configurations of Chassis Wide Settings

Router A

```
user@host> show configuration forwarding-options hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric hash;
}
```

Router B

```
user@host> show configuration forwarding-options hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric-hash {
    complement;
  }
}
```

Example Configurations of Per–Packet-Forwarding-Engine Settings

Router A

```
user@host> show configuration chassis fpc 2 pic 2 hash-key
family multiservice {
  payload {
    ip {
      layer-3;
    }
  }
  symmetric hash;
}
```

Copyright © 2013, Juniper Networks, Inc.

Router B

```
user@host> show configuration chassis fpc 2 pic 3 hash-key
family multiservice {
   payload {
      ip {
         layer-3;
      }
   }
   symmetric-hash {
      complement;
   }
}
```

Related Documentation

- *Ethernet Interfaces*
- For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library for Routing Devices*.

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 5

Stacking and Rewriting Gigabit Ethernet VLAN Tags

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

Stacking and rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between customer edge (CE) routers that share one VLAN ID. A frame can be received on an interface, or it can be internal to the system (as a result of the **input-vlan-map** statement).

On IQ2 interfaces, 10-Gigabit Ethernet LAN/WAN PIC, 40-Gigabit Ethernet MIC, 100-Gigabit Ethernet MIC, IQ2-E interfaces, and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

Copyright © 2013, Juniper Networks, Inc.

NOTE: When swap-by-poppush is configured on the interface, when a VLAN tag is swapped, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN being swapped. If swap-by-poppush is not configured on the interface, the VLAN IEEE 802.1p bits of the of the VLAN being swapped remains same.

You can stack and rewrite VLAN tags on the following interfaces:

- Gigabit Ethernet
- Gigabit Ethernet IQ
- 10-Gigabit Ethernet LAN/WAN PIC
- 40-Gigabit Ethernet MIC
- 100-Gigabit Ethernet MIC
- Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, and MX Series router Gigabit Ethernet Interfaces
- Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with the VLAN encapsulation type configured to support Layer 2 tunneling protocols such as circuit cross-connect (CCC) or virtual private LAN service (VPLS) (as described in "802.1Q VLANs Overview" on page 53)

Related Documentation

- 802.1Q VLANs Overview on page 53
- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 198

- *Ethernet Interfaces*

Stacking and Rewriting Gigabit Ethernet VLAN Tags

You can configure rewrite operations to stack (**push**), remove (**pop**), or rewrite (**swap**) tags on single-tagged frames and dual-tagged frames. If a port is not tagged, rewrite operations are not supported on any logical interface on that port.

You can configure the following VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **pop-pop**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.

 Copyright © 2013, Juniper Networks, Inc.

- **push**—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- **push-push**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, push two VLAN tags in front of the frame.
- **swap-push**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.
- **swap-swap**—For Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces, replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value.

You configure VLAN rewrite operations for logical interfaces in the input VLAN map for incoming frames and in the output VLAN map for outgoing frames. To configure the input VLAN map, include the **input-vlan-map** statement:

```
input-vlan-map {
  ...interface-specific configuration...
}
```

To configure the output VLAN map, include the **output-vlan-map** statement:

```
output-vlan-map {
  ...interface-specific configuration...
}
```

You can include both statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The type of VLAN rewrite operation permitted depends upon whether the frame is single-tagged or dual-tagged. Table 9 on page 199 shows supported rewrite operations and whether they can be applied to single-tagged frames or dual-tagged frames. The table also indicates the number of tags being added or removed during the operation.

Table 9: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames

Rewrite Operation	Untagged	Single-Tagged	Dual-Tagged	Number of Tags
pop	No	Yes	Yes	– 1
push	Sometimes	Yes	Yes	+1
swap	No	Yes	Yes	0
push-push	Sometimes	Yes	Yes	+2
swap-push	No	Yes	Yes	+1

Copyright © 2013, Juniper Networks, Inc.

Table 9: Rewrite Operations on Untagged, Single-Tagged, and Dual-Tagged Frames *(continued)*

Rewrite Operation	Untagged	Single-Tagged	Dual-Tagged	Number of Tags
swap-swap	No	No	Yes	0
pop-pop	No	No	Yes	– 2
pop-swap	No	No	Yes	– 1

The rewrite operations **push** and **push-push** can be valid in certain circumstances on frames that are not tagged. For example, a single-tagged logical interface (interface 1) and a dual-tagged logical interface (interface 2) have the following configurations:

Interface 1

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
   pop;
}
output-vlan-map {
   push;
}
```

Interface 2

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
   pop-pop;
}
output-vlan-map {
   push-push;
}
```

When a frame is received on the interface as a result of the **input-vlan-map** operation, the frame is not tagged. As it goes out of the second interface, the **output-vlan-map** operation **push-push** is applied to it. The resulting frame will be dual-tagged at the logical interface output.

Depending on the VLAN rewrite operation, you configure the rewrite operation for the interface in the input VLAN map, the output VLAN map, or in both the input VLAN map and the output VLAN map. Table 10 on page 200 shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

Table 10: Applying Rewrite Operations to VLAN Maps

	Output VLAN Map								
Input VLAN Map	none	push	pop	swap	push-push	swap-push	swap-swap	pop-pop	swap-pop
none	Yes	No	No	Yes	No	No	Yes	No	No
push	No	No	Yes	No	No	No	No	No	No
pop	No	Yes	No	No	No	No	No	No	No

Copyright © 2013, Juniper Networks, Inc.

Table 10: Applying Rewrite Operations to VLAN Maps *(continued)*

	Output VLAN Map								
Input VLAN Map	none	push	pop	swap	push-push	swap-push	swap-swap	pop-pop	swap-pop
swap	Yes	No	No	Yes	No	No	No	No	No
push-push	No	No	No	No	No	No	No	Yes	No
swap-push	No	No	No	No	No	No	No	No	Yes
swap-swap	Yes	No	No	No	No	No	Yes	No	No
pop-pop	No	No	No	No	Yes	No	No	No	No
pop-swap	No	No	No	No	No	Yes	No	No	No

You must know whether the VLAN rewrite operation is valid and is applied to the input VLAN map or the output VLAN map. You must also know whether the rewrite operation requires you to include statements to configure the inner and outer TPIDs and inner and outer VLAN IDs in the input VLAN map or output VLAN map. For information about configuring inner and outer TPIDs and inner and outer VLAN IDs, see "Configuring Inner and Outer TPIDs and VLAN IDs" on page 203.

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 197
- *Understanding swap-by-poppush*
- *swap-by-poppush*
- *Ethernet Interfaces*

Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames

For Gigabit Ethernet IQ interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure frames with particular TPIDs to be processed as tagged frames. To do this, you specify up to eight IEEE 802.1Q TPID values per port; a frame with any of the specified TPIDs is processed as a tagged frame; however, with IQ2 and IQ2-E interfaces, only the first four IEEE 802.1Q TPID values per port are supported. To configure the TPID values, include the **tag-protocol-id** statement:

```
tag-protocol-id [ tpids ];
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]**
- **[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]**

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** or **[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

Related Documentation

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 197
- **aggregated-ether-options on page 555**
- **ethernet-switch-profile on page 568**
- **gigether-options on page 577**
- **tag-protocol-id on page 659**
- *Ethernet Interfaces*

Configuring Stacked VLAN Tagging

To configure stacked VLAN tagging for all logical interfaces on a physical interface, include the **stacked-vlan-tagging** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
stacked-vlan-tagging;
```

If you include the **stacked-vlan-tagging** statement in the configuration, you must configure dual VLAN tags for all logical interfaces on the physical interface. For more information, see "Stacking a VLAN Tag" on page 206.

Related Documentation

- *stacked-vlan-tagging*
- Stacking a VLAN Tag on page 206
- *Ethernet Interfaces*

Configuring Dual VLAN Tags

To configure dual VLAN tags on a logical interface, include the **vlan-tags** statement:

```
vlan-tags inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The outer tag VLAN ID range is from 1 through 511 for normal interfaces, and from 512 through 4094 for VLAN CCC or VLAN VPLS interfaces. The inner tag is not restricted.

You must also include the **stacked-vlan-tagging** statement in the configuration. See "Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags" on page 215.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- unit on page 667
- Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 215
- *Ethernet Interfaces*

Configuring Inner and Outer TPIDs and VLAN IDs

For some rewrite operations, you must configure the inner or outer TPID values and inner or outer VLAN ID values. These values can be applied to either the input VLAN map or the output VLAN map.

On Ethernet IQ, IQ2, and IQ2-E interfaces; on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces; and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to configure the inner TPID, include the **inner-tag-protocol-id** statement:

```
inner-tag-protocol-id tpid;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

For the inner VLAN ID, include the **inner-vlan-id** statement. For the outer TPID, include the **tag-protocol-id** statement. For the outer VLAN ID, include the **vlan-id** statement:

```
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;
}
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;
}
```

For aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, include the **tag-protocol-id** statement for the outer TPID. For the outer VLAN ID, include the **vlan-id** statement:

```
input-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
```

Copyright © 2013, Juniper Networks, Inc.

```
    vlan-id number;
  }
  output-vlan-map {
    (pop | push | swap);
    tag-protocol-id tpid;
    vlan-id number;
  }
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see "802.1Q VLANs Overview" on page 53.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level or **[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level. For more information, see "Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames" on page 201.

Table 11 on page 204 and Table 12 on page 205 specify when these statements are required. Table 11 on page 204 indicates valid statement combinations for rewrite operations for the input VLAN map. "No" means the statement must not be included in the input VLAN map for the rewrite operation. "Optional" means the statement may be optionally specified for the rewrite operation in the input VLAN map. "Any" means that you must include the **vlan-id** statement, **tag-protocol-id** statement, **inner-vlan-id** statement, or **inner-tag-protocol-id** statement.

Table 11: Rewrite Operations and Statement Usage for Input VLAN Maps

	Input VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Optional	Optional	No	No
pop	No	No	No	No
swap	Any	Any	No	No
push-push	Optional	Optional	Optional	optional
swap-push	Optional	Optional	Any	Any
swap-swap	Optional	Optional	Any	Any
pop-swap	No	No	Any	Any

 Copyright © 2013, Juniper Networks, Inc.

Table 11: Rewrite Operations and Statement Usage for Input VLAN Maps *(continued)*

	Input VLAN Map Statements			
pop-pop	No	No	No	No

Table 12 on page 205 indicates valid statement combinations for rewrite operations for the output VLAN map. “No” means the statement must not be included in the output VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 12: Rewrite Operations and Statement Usage for Output VLAN Maps

	Output VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	No	Optional	No	No
pop	No	No	No	No
swap	No	Optional	No	No
push-push	No	Optional	No	Optional
swap-push	No	Optional	No	Optional
swap-swap	No	Optional	No	Optional
pop-swap	No	No	No	Optional
pop-pop	No	No	No	No

The following examples use Table 11 on page 204 and Table 12 on page 205 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

Input VLAN Map with inner-vlan-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id Statement, Output

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
```

Copyright © 2013, Juniper Networks, Inc.

VLAN Map with Optional inner-tag-protocol-id Statement

```
  inner-tag-protocol-id tpid;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
  inner-tag-protocol-id tpid;
}
```

Related Documentation

- inner-tag-protocol-id on page 581
- input-vlan-map on page 585
- *output-vlan-map*
- pop-swap on page 624
- unit on page 667
- *Ethernet Interfaces*

Stacking a VLAN Tag

To stack a VLAN tag on all tagged frames entering or exiting the interface, include the **push**, **vlan-id**, and **tag-protocol-id** statements in the input VLAN map or the output VLAN map:

```
input-vlan-map {
  push;
  vlan-id number;
  tag-protocol-id tpid;
}
output-vlan-map {
  push;
  tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

If you include the **push** statement in an interface's input VLAN map, see Table 10 on page 200 for information about permissible rewrite operations,

 Copyright © 2013, Juniper Networks, Inc.

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see "802.1Q VLANs Overview" on page 53.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level. For more information, see "Configuring Inner and Outer TPIDs and VLAN IDs" on page 203.

Related Documentation

- **tag-protocol-id on page 660**
- **unit on page 667**
- Table 10 on page 200
- 802.1Q VLANs Overview on page 53
- Configuring Inner and Outer TPIDs and VLAN IDs on page 203
- *Ethernet Interfaces*

Removing a VLAN Tag

To remove a VLAN tag from all tagged frames entering or exiting the interface, include the **pop** statement in the input VLAN map or output VLAN map:

```
pop;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

Related Documentation

- **input-vlan-map on page 585**
- *output-vlan-map*
- **pop on page 623**
- **unit on page 667**
- *Ethernet Interfaces*

Removing the Outer and Inner VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series

Copyright © 2013, Juniper Networks, Inc.

routers, to remove both the outer and inner VLAN tags of the frame, include the **pop-pop** statement in the input VLAN map or output VLAN map:

```
pop-pop;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- **input-vlan-map on page 585**
- *output-vlan-map*
- **pop-pop on page 623**
- **unit on page 667**
- See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- *Ethernet Interfaces*

Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame with a user-specified VLAN tag value, include the **pop-swap** statement in the input VLAN map or output VLAN map:

```
pop-swap;
```

The inner tag becomes the outer tag in the final frame.

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

 Copyright © 2013, Juniper Networks, Inc.

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 585
- *output-vlan-map*
- pop-swap on page 624
- unit on page 667
- See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- *Ethernet Interfaces*

Stacking Two VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to push two VLAN tags in front of tagged frames entering or exiting the interface, include the **push-push** statement in the input VLAN map or the output VLAN map:

```
push-push;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 585
- *output-vlan-map*
- pop on page 623
- unit on page 667

- See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.
- *Ethernet Interfaces*

Rewriting the VLAN Tag on Tagged Frames

To rewrite the VLAN tag on all tagged frames entering the interface to a specified VLAN ID and TPID, include the **swap**, **tag-protocol-id**, and **vlan-id** statements in the input VLAN map:

```
input-vlan-map {
  swap;
  vlan-id number;
  tag-protocol-id tpid;
}
```

To rewrite the VLAN tag on all tagged frames exiting the interface to a specified VLAN ID and TPID, include the **swap** and **tag-protocol-id** statements in the output VLAN map:

```
output-vlan-map {
  swap;
  vlan-id number;
  tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**

You cannot include both the **swap** statement and the **vlan-id** statement in the output VLAN map configuration. If you include the **swap** statement in the configuration, the VLAN ID in outgoing frames is rewritten to the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see “802.1Q VLANs Overview” on page 53.

The swap operation works on the outer tag only, whether or not you include the **stacked-vlan-tagging** statement in the configuration. For more information, see “Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags” on page 215.

Related Documentation

- input-vlan-map on page 585
- *output-vlan-map*
- swap on page 653
- vlan-id on page 676
- tag-protocol-id on page 660
- unit on page 667

 Copyright © 2013, Juniper Networks, Inc.

- For more information about binding a VLAN ID to the logical interface, see 802.1Q VLANs Overview on page 53.
- For more information about the swap operation, see Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 215.
- *Ethernet Interfaces*

Rewriting a VLAN Tag on Untagged Frames

On M320, M120, and MX Series routers with Gigabit Ethernet IQ, IQ2, and IQ2E PICs, 10-Gigabit Ethernet IQ, IQ2, and IQ2E PICs, and on MX Series 40-port Gigabit Ethernet R, 40-port Gigabit Ethernet R EQ, 4-port 10-Gigabit Ethernet R, and 4-port 10-Gigabit Ethernet R EQ DPCs, you can rewrite VLAN tags on untagged incoming and outgoing frames under **ethernet-ccc** and **ethernet-vpls** encapsulations. On MX Series routers with IQ2 and IQ2-E PICs, you can perform all rewrite VLAN tag operations. These features provide added flexibility.

Consider a network where two provider edges (PE) are connected by a Layer 2 circuit. PE1 is receiving traffic on an untagged port while the corresponding port on PE2 is tagged. In the normal case, packets coming from PE1 will be dropped at PE2 because it is expecting tagged packets. However, if PE1 can push a VLAN tag on the incoming packet before sending it across to PE2, you can ensure that packets are not dropped. To make it work in both directions, PE1 must strip the VLAN tag from outgoing packets. Therefore, a push on the ingress side is always paired with a pop on the egress side.

The rewrite operations represented by the following statement options are supported under **ethernet-ccc** and **ethernet-vpls** encapsulations:

- **push**—A VLAN tag is added to the incoming untagged frame.
- **pop**—VLAN tag is removed from the outgoing frame.
- **push-push**—An outer and inner VLAN tag are added to the incoming untagged frame.
- **pop-pop**—Both the outer and inner VLAN tags of the outgoing frame are removed.

IQ2 and 10-Gigabit Ethernet PICs support all rewrite operations described above. Details on the possible combinations of usage are explained later in this section.

NOTE: **The push-push and pop-pop operations are not supported on the Gigabit Ethernet IQ PIC.**

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the tag parameters have to be explicitly specified. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames. Table 13 on page 212 through Table 15 on page 212 explain the rules in more detail.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the **vlan-id** parameters (**vlan-id** for **push** and **vlan-id** or **inner-vlan-id** for **push-push**) have to be explicitly specified. TPID however, is optional and the default value of **0x8100** is set if not configured. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames.

Table 13: Input VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Yes	Optional	No	Optional
push-push	Yes	Optional	Yes	Optional

Table 14: Output VLAN Map Statements Allowed for ethernet-ccc and ethernet-vpls Encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
pop	No	No	No	No
pop-pop	No	No	No	No

Table 15: Rules for Applying Rewrite Operations to VLAN Maps

	Output VLAN Map		
Input VLAN Map	None	pop	pop-pop
None	Yes	No	No
push	No	Yes	No
push-push	No	No	Yes

Example: push and pop with Ethernet CCC Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 600;
    }
    output-vlan-map pop;
    family ccc;
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

Example: push-push and pop-pop with Ethernet CCC Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family ccc;
  }
}
```

Example: push and pop with Ethernet VPLS Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 700;
    }
    output-vlan-map pop;
    family vpls;
  }
}
```

Example: push-push and pop-pop with Ethernet VPLS Encapsulation

```
ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family vpls;
  }
}
```

You can use the **show interface** ***interface-name*** command to display the status of a modified VLAN map for the specified interface.

Related Documentation

- input-vlan-map on page 585
- *output-vlan-map*
- pop on page 623
- pop-pop on page 623

- push on page 635
- push-push on page 636
- unit on page 667
- *Ethernet Interfaces*

Rewriting a VLAN Tag and Adding a New Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and on Gigabit Ethernet and 10-Gigabit Ethernet interfaces on EX Series switches, to replace the outer VLAN tag of the incoming frame with a user-specified VLAN tag value, include the **swap-push** statement in the input VLAN map or output VLAN map:

```
swap-push
```

A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- input-vlan-map on page 585
- *output-vlan-map*
- swap-push on page 654
- unit on page 667
- *Ethernet Interfaces*

Rewriting the Inner and Outer VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to replace both the inner and the outer VLAN tags of the incoming frame with a

 Copyright © 2013, Juniper Networks, Inc.

user-specified VLAN tag value, include the **swap-swap** statement in the input VLAN map or output VLAN map:

```
swap-swap;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

See Table 11 on page 204 and Table 12 on page 205 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- **input-vlan-map on page 585**
- *output-vlan-map*
- **swap-swap on page 655**
- **unit on page 667**
- *Ethernet Interfaces*

Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags

Configure a VLAN CCC tunnel in which Ethernet frames enter the tunnel at interface **ge-4/0/0** and exit the tunnel at interface **ge-4/2/0**.

The following examples show how to perform the following tasks:

- Push a TPID and VLAN ID Pair on Ingress on page 216
- Stack Inner and Outer VLAN Tags on page 217
- Swap a VLAN ID on Ingress on page 217
- Swap a VLAN ID on Egress on page 218
- Swap a VLAN ID on Both Ingress and Egress on page 218
- Swap the Outer VLAN Tag and Push a New VLAN Tag on Ingress; Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Egress on page 219
- Swap a TPID and VLAN ID Pair for Both VLAN Tags on Ingress and on Egress on page 220
- Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress on page 220

- Pop a TPID and VLAN ID Pair on Ingress; Push a VLAN ID and TPID Pair on Egress on page 220
- Pop an Outer VLAN Tag to Connect an Untagged VPLS Interface to Tagged VPLS Interfaces on page 221

Push a TPID and VLAN ID Pair on Ingress

```
[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  gigether-options {
    ethernet-switch-profile {
      tag-protocol-id 0x9909;
    }
  }
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 512;
    input-vlan-map {
      push;
      tag-protocol-id 0x9909;
      vlan-id 520;
    }
    output-vlan-map pop;
  }
}
ge-4/2/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 515;
    input-vlan-map {
      swap-push;
      vlan-id 520;
      inner-vlan-id 512;
    }
    output-vlan-map {
      pop-swap;
    }
  }
  [edit protocols]
  mpls {
    interface ge-4/0/0.0;
    interface ge-4/2/0.0;
  }
  connections {
    interface-switch vlan-tag-push {
      interface ge-4/0/0.0;
      interface ge-4/2/0.0;
    }
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

Stack Inner and Outer VLAN Tags

```
[edit interfaces]
ge-0/2/0 {
  stacked-vlan-tagging;
  mac 00.01.02.03.04.05;
  gigether-options {
    loopback;
  }
  unit 0 {
    vlan-tags outer 0x8100.200 inner 0x8100.200;
  }
}
```

Swap a VLAN ID on Ingress

```
[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  gigether-options {
    ethernet-switch-profile {
      tag-protocol-id 0x9100;
    }
  }
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
    input-vlan-map {
      swap;
      tag-protocol-id 0x9100;
      vlan-id 2000;
    }
  }
}
ge-4/2/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 2000;
    input-vlan-map {
      swap;
      tag-protocol-id 0x9100;
      vlan-id 1000;
    }
  }
  [edit protocols]
  mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
  }
  connections {
    ...
    interface-switch vlan-tag-swap {
      interface ge-4/2/0.1;
```

```
      interface ge-4/0/0.1;
    }
  }
}
```

Swap a VLAN ID on Egress

```
[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
  }
}
ge-4/2/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  gigether-options {
    ethernet-switch-profile {
      tag-protocol-id 0x8800;
    }
  }
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 2000;
    output-vlan-map {
      swap;
      tag-protocol-id 0x8800;
    }
  }
}
[edit protocols]
mpls {
  ...
  interface ge-4/0/0.1;
  interface ge-4/2/0.1;
}
connections {
  ...
  interface-switch vlan-tag-swap {
    interface ge-4/2/0.1;
    interface ge-4/0/0.1;
  }
}
```

Swap a VLAN ID on Both Ingress and Egress

```
[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  gigether-options {
    ethernet-switch-profile {
      tag-protocol-id [ 0x8800 0x9100 ];
    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
    ...
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 1000;
        input-vlan-map {
            swap;
            tag-protocol-id 0x9100;
            vlan-id 2000;
        }
    }
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        output-vlan-map {
            swap;
            tag-protocol-id 0x8800;
        }
    }
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}
```

Swap the Outer VLAN Tag and Push a New VLAN Tag on Ingress; Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Egress

```
[edit interfaces]
ge-1/1/0 {
    unit 1 {
        vlan-id 200;
        input-vlan-map {
            swap-push;
            tag-protocol-id 0x9100;
            vlan-id 400;
            inner-tag-protocol-id 0x9100;
            inner-vlan-id 500;
        }
        output-vlan-map {
            pop-swap;
```

Copyright © 2013, Juniper Networks, Inc.

```
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

Swap a TPID and VLAN ID Pair for Both VLAN Tags on Ingress and on Egress

```
[edit interfaces]
ge-1/1/0 {
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map {
      swap-swap;
      tag-protocol-id 0x9100;
      vlan-id 400;
      inner-tag-protocol-id 0x9100;
      inner-vlan-id 500;
    }
    output-vlan-map {
      swap-swap;
      tag-protocol-id 0x9200;
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress

```
[edit interfaces]
ge-1/1/0 {
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map {
      pop-swap;
      tag-protocol-id 0x9100;
      vlan-id 400;
    }
    output-vlan-map{
      swap-push;
      tag-protocol-id 0x9200;
      inner-tag-protocol-id 0x9100;
    }
  }
}
```

Pop a TPID and VLAN ID Pair on Ingress; Push a VLAN ID and TPID Pair on Egress

```
[edit interfaces]
ge-1/1/0 {
  unit 0 {
    vlan-tags {
      inner 0x9100.425;
      outer 0x9200.525;
    }
    input-vlan-map{
```

 Copyright © 2013, Juniper Networks, Inc.

```
            pop-pop;
        }
        output-vlan-map {
            push-push;
            tag-protocol-id 0x9200;
            inner-tag-protocol-id 0x9100;
        }
    }
}
```

Pop an Outer VLAN Tag to Connect an Untagged VPLS Interface to Tagged VPLS Interfaces

```
[edit interfaces]
ge-1/1/0 {
    vlan-tagging;
    encapsulation extended-vlan-vpls;
    unit 0 {
        vlan-id 0;
        input-vlan-map {
            push;
            vlan-id 0;
        }
        output-vlan-map pop;
        family vpls;
    }
}
```

Related Documentation

- input-vlan-map on page 585
- *output-vlan-map*
- inner-tag-protocol-id on page 581
- inner-vlan-id on page 582
- pop on page 623
- pop-pop on page 623
- pop-swap on page 624
- push on page 635
- push-push on page 636
- swap on page 653
- swap-push on page 654
- swap-swap on page 655
- unit on page 667
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 6

Configuring Layer 2 Bridging Interfaces

Layer 2 Bridging Interfaces Overview

Bridging operates at Layer 2 of the OSI reference model while routing operates at Layer 3. A set of logical ports configured for bridging can be said to constitute a bridging domain.

A bridging domain can be created by configuring a routing instance and specifying the instance-type as **bridge**.

Integrated routing and bridging (IRB) is the ability to:

- Route a packet if the destination MAC address is the MAC address of the router and the packet **ethertype** is IPv4, IPv6, or MPLS.
- Switch all multicast and broadcast packets within a bridging domain at layer 2.
- Route a copy of the packet if the destination MAC address is a multicast address and the **ethertype** is IPv4 or IPv6.
- Switch all other unicast packets at Layer 2.
- Handle supported Layer 2 control packets such as STP and LACP.
- Handle supported Layer 3 control packets such as OSPF and RIP.

Related Documentation

- Configuring Layer 2 Bridging Interfaces on page 223
- *Ethernet Interfaces*

Configuring Layer 2 Bridging Interfaces

You can configure an IRB logical interface at the **[edit interfaces ge-*fpc* /*pic*/*port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
unit logical-unit-number {
}
```

You can configure Layer 3 information on the IRB logical interface by including the **irb** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
irb {
   unit logical-unit-number {
      family inet {
         address address {
         }
      }
   }
}
```

For examples of Layer 2 bridging configuration, see the *Junos OS Routing Protocols Library for Routing Devices*.

Example: Configuring Layer 2 Bridging Interfaces

The following example configures an IRB logical interface and Layer 3 information on the interface.

```
[edit interfaces]
ge-1/0/0 {
   unit 0 {
   }
}
irb {
   unit 0 {
      family inet {
      address 192.168.12.1/28;
      }
   }
}
```

Related Documentation

- *family*
- **unit on page 667**
- Layer 2 Bridging Interfaces Overview on page 223
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 7

Configuring Link Layer Discovery Protocol

LLDP Overview

The Link Layer Discovery Protocol (LLDP) is an industry-standard, vendor-neutral method to allow networked devices to advertise capabilities, identity, and other information onto a LAN. The Layer 2 protocol, detailed in IEEE 802.1AB-2005, replaces several proprietary protocols implemented by individual vendors for their equipment.

LLDP allows network devices that operate at the lower layers of a protocol stack (such as Layer 2 bridges and switches) to learn some of the capabilities and characteristics of LAN devices available to higher layer protocols, such as IP addresses. The information gathered through LLDP operation is stored in a network device and is queried with SNMP. Topology information can also be gathered from this database.

Some of the information that can be gathered by LLDP (only minimal information is mandatory) is:

- System name and description
- Port name and description
- VLAN name and identifier
- IP network management address
- Capabilities of the device (for example, switch, router, or server)
- MAC address and physical layer information
- Power information
- Link aggregation information

NOTE: LLDP media endpoint discovery (LLDP-MED) is not supported on T Series routers.

LLDP frames are sent at fixed intervals on each port that runs LLDP. LLDP protocol data units (LLDP PDUs) are sent inside Ethernet frames and identified by their destination Media Access Control (MAC) address (**01:80:C2:00:00:0E**) and Ethertype (**0x88CC**). Mandatory information supplied by LLDP is chassis ID, port ID, and a time-to-live value for this information.

LLDP is a powerful way to allow Layer 2 devices to gather details about other network-attached devices.

Related Documentation

- Configuring LLDP on page 226
- Tracing LLDP Operations on page 228
- Example: Configuring LLDP on page 229

- *LLDP Operational Mode Commands*

Configuring LLDP

You configure LLDP by including the **lldp** statement and associated parameters at the **[edit protocols]** hierarchy level. The complete set of LLDP statements follows:

```
lldp {
  advertisement-interval seconds;
  disable;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  port-id-subtype {
    interface-name;
    locally-assigned;
  }
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  transmit-delay seconds
}
```

The following statements have default values:

- **advertisement-interval**—The default value is 30 seconds. The allowable range is from 5 through 32768 seconds.
- **hold-multiplier**—The default values is 4. The allowable range is from 2 through 10.
- **ptopo-configuration-maximum-hold-time**—The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds.
- **transmit-delay**—The default values is 2 seconds. The allowable range is from 1 through 8192 seconds.

Copyright © 2013, Juniper Networks, Inc.

The following statements must be explicitly configured:

- **lldp-configuration-notification-interval**—The allowable range is from 0 through 3600 seconds. There is no default value.
- **ptopo-configuration-trap-interval**—The allowable range is from 1 through 2147483647 seconds. There is no default value.

To disable LLDP on all or a particular interface, include the **interfaces** statement at the **[edit protocols lldp]** hierarchy level:

```
interface (all | interface-name) {
  disable;
}
```

To disable LLDP on all interfaces, use the **all** option. To disable LLDP on a particular interface, include the **disable** statement with the interface name.

To configure LLDP on a T Series router within a TX Matrix, you must specify the interface name in the LLDP configuration for the TX Matrix. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering, see *Routing Matrix with a TX Matrix Router FPC Numbering*

NOTE: **The interface-name must be the physical interface (for example, ge-1/0/0) and not a logical interface (unit).**

The advertisement interval determines the frequency that an LLDP interface sends LLDP advertisement frames. The default value is 30 seconds. The allowable range is from 5 through 32768 seconds. You adjust this parameter by including the **advertisement-interval** statement at the **[edit protocols lldp]** hierarchy level.

The hold multiplier determines the multiplier to apply to the advertisement interval. The resulting value in seconds is used to cache learned LLDP information before discard. The default value is 4. When used with the default advertisement interval value of 30 seconds, this makes the default cache lifetime 120 seconds. The allowable range of the hold multiplier is from 2 through 10. You adjust this parameter by including the **hold-multiplier** statement at the **[edit protocols lldp]** hierarchy level.

The transmit delay determines the delay between any two consecutive LLDP advertisement frames. The default value is 2 seconds. The allowable range is from 1 through 8192 seconds. You adjust this parameter by including the **transmit-delay** statement at the **[edit protocols lldp]** hierarchy level.

The physical topology configuration maximum hold time determines the time interval for which an agent device maintains physical topology database entries. The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds. You adjust this parameter by including the **ptopo-configuration-maximum-hold-time** statement at the **[edit protocols lldp]** hierarchy level.

The LLDP configuration notification interval determines the period for which trap notifications are sent to the SNMP Master Agent when changes occur in the database

Copyright © 2013, Juniper Networks, Inc.

of LLDP information. This capability is disabled by default. The allowable range is from 0 (disabled) through 3600 seconds. You adjust this parameter by including the **lldp-configuration-notification-interval** statement at the **[edit protocols lldp]** hierarchy level.

The physical topology configuration trap interval determines the period for which trap notifications are sent to the SNMP Master Agent when changes occur in the global physical topology statistics. This capability is disabled by default. The allowable range is from 0 (disabled) through 3600 seconds. The LLDP agent sends traps to the SNMP Master Agent if this interval has a value greater than 0 and there is any change during the **lldp-configuration-notification-interval** trap interval. You adjust this parameter by including the **ptopo-configuration-trap-interval** statement at the **[edit protocols lldp]** hierarchy level.

By default, LLDP generates the SNMP index of the interface for the port ID Type, Length, and Value (TLV). Starting with Junos OS Release 12.3R1, you can generate the interface name as the port ID TLV by including the **interface-name** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level. When configure the **interface-name**, statement on the remote LLDP neighbor, the **show lldp neighbors** command displays the interface name in the **Port ID** field rather than the SNMP index of the interface, which is displayed by default. If you change the default behavior of generating the SNMP index of the interface as the Port ID TLV, you can reenable the default behavior by including the **locally-assigned** statement at the **[edit protocols lldp port-id-subtype** hierarchy level.

Related Documentation

- LLDP Overview on page 225
- Tracing LLDP Operations on page 228
- Example: Configuring LLDP on page 229
- *LLDP Operational Mode Commands*
- *TX Matrix Router Chassis and Interface Names*
- *Miscellaneous Commands for a Routing Matrix with a TX Matrix Router*

Tracing LLDP Operations

To trace LLDP operational traffic, you can specify options in the global **traceoptions** statement included at the **[edit routing-options]** hierarchy level, and you can specify LLDP-specific options by including the **traceoptions** statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols lldp]**
- **[edit routing-instances *routing-instance-name* protocols lldp]**

 Copyright © 2013, Juniper Networks, Inc.

You can specify the following LLDP-specific options in the LLDP **traceoptions** statement:

- **all**—Trace all operations.
- **config**—Log configuration events.
- **interface**—Trace interface update events.
- **protocol**—Trace protocol information.
- **rtsock**—Trace real-time socket events.
- **vlan**—Trace VLAN update events.

NOTE: Use the trace flag **all** with caution. This flag may cause the CPU to become very busy.

For general information about tracing and global tracing options, see the statement summary for the global **traceoptions** statement in the *Junos OS Routing Protocols Library for Routing Devices*.

Related Documentation

- LLDP Overview on page 225
- Configuring LLDP on page 226
- Example: Configuring LLDP on page 229
- *LLDP Operational Mode Commands*

Example: Configuring LLDP

The following example configures LLDP on interface **ge-1/1/1** but disables LLDP on all other interfaces, explicitly configures the default values for all automatically enabled features, and configures a value of 30 seconds for the LLDP configuration notification interval and a value of 30 seconds for the physical topology trap interval.

```
[edit]
protocols {
   lldp {
      advertisement-interval 30;
      hold-multiplier 4;
      interface all {
         disable;
      }
      interface ge-1/1/1;
      lldp-configuration-notification-interval 30;
      ptopo-configuration-maximum-hold-time 300;
      ptopo-configuration-trap-interval 30;
      transmit-delay 2;
   }
}
```

Copyright © 2013, Juniper Networks, Inc.

You verify operation of LLDP with several show commands:

- **show lldp <detail>**
- **show lldp neighbors** ***interface-name***
- **show lldp statistics** ***interface-name***
- **show lldp local-information**
- **show lldp remote-global-statisitics**

You can clear LLDP neighbor information or statistics globally or on an interface:

- **clear lldp neighbors** ***interface-name***
- **clear lldp statistics** ***interface-name***

You can display basic information about LLDP with the **show lldp detail** command:

```
user@host> show lldp detail
LLDP                    : Enabled
Advertisement interval  : 30 Second(s)
Transmit delay          : 2 Second(s)
Hold timer              : 4 Second(s)
Notification interval   : 30 Second(s)
Config Trap Interval    : 300 Second(s)
Connection Hold timer   : 60 Second(s)


Interface      LLDP       Neighbor count
ge-1/1/1       Enabled    0

LLDP basic TLVs supported:
Chassis identifier, Port identifier, Port description, System name, System
description, System capabilities, Management address.

LLDP 802 TLVs supported:
Link aggregation, Maximum frame size, MAC/PHY Configuration/Status, Port VLAN ID,
 Port VLAN name.
```

For more details about the output of these commands, see the *Junos OS Operational Mode Commands*.

Related Documentation

- LLDP Overview on page 225
- Configuring LLDP on page 226
- Tracing LLDP Operations on page 228
- *LLDP Operational Mode Commands*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 8

Configuring TCC and Layer 2.5 Switching

TCC and Layer 2.5 Switching Overview

Translational cross-connect (TCC) is a switching concept that allows you to forward traffic between a variety of Layer 2 protocols or circuits. It is similar to its predecessor, CCC. However, while CCC requires the same Layer 2 encapsulations on both sides of a router (such as Point-to-Point Protocol [PPP] or Frame Relay-to-Frame Relay), TCC lets you connect different types of Layer 2 protocols interchangeably. With TCC, combinations such as PPP-to-ATM and Ethernet-to-Frame Relay cross-connections are possible.

Related Documentation

- Configuring VLAN TCC Encapsulation on page 231
- Configuring Ethernet TCC on page 233
- *Ethernet Interfaces*

Configuring VLAN TCC Encapsulation

VLAN TCC encapsulation allows circuits to have different media on either side of the forwarding path. VLAN TCC encapsulation supports TPID 0x8100 only. You must include configuration statements at the logical and physical interface hierarchy levels.

To configure VLAN TCC encapsulation, include the **encapsulation** statement and specify the **vlan-tcc** option:

[edit interfaces *interface-name* unit *logical-unit-number*]
encapsulation vlan-tcc;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Copyright © 2013, Juniper Networks, Inc.

Additionally, configure the logical interface by including the **proxy** and **remote** statements:

```
proxy {
  inet-address;
}
remote {
  (inet-address | mac-address);
}
```

You can include these statements at the following hierarchy levels:

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

The proxy address is the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy.

The remote address is the IP or MAC address of the remote router. The **remote** statement provides ARP capability from the TCC switching router to the Ethernet neighbor. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

When VLAN TCC encapsulation is configured on the logical interface, you also must specify flexible Ethernet services on the physical interface. To specify flexible Ethernet services, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level and specify the **flexible-ethernet-services** option:

```
[edit interfaces interface-name]
encapsulation flexible-ethernet-services;
```

Extended VLAN TCC encapsulation supports TPIDs 0x8100 and 0x9901. Extended VLAN TCC is specified at the physical interface level. When configured, all units on that interface must use VLAN TCC encapsulation, and no explicit configuration is needed on logical interfaces.

One-port Gigabit Ethernet, 2-port Gigabit Ethernet, and 4-port Fast Ethernet PICs with VLAN tagging enabled can use VLAN TCC encapsulation. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level and specify the **extended-vlan-tcc** option:

```
[edit interfaces interface-name]
encapsulation extended-vlan-tcc;
```

For VLAN TCC encapsulation, all VLAN IDs from 1 through 1024 are valid. VLAN ID 0 is reserved for tagging the priority of frames.

Extended VLAN TCC is not supported on 4-port Gigabit Ethernet PICs.

Related Documentation

- *encapsulation*
- *remote*
- *proxy*
- TCC and Layer 2.5 Switching Overview on page 231

Copyright © 2013, Juniper Networks, Inc.

- Configuring Ethernet TCC on page 233
- *Ethernet Interfaces*

Configuring Ethernet TCC

For Layer 2.5 virtual private networks (VPNs) using an Ethernet interface as the TCC router, you can configure an Ethernet TCC.

To configure an Ethernet TCC, include the **encapsulation** statement and specify the **ethernet-tcc** option at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation ethernet-tcc;
```

For Ethernet TCC encapsulation, you must also configure the logical interface by including the **proxy** and **remote** statements:

```
proxy {
  inet-address;
}
remote {
  (inet-address | mac-address);
}
```

You can include these statements at the following hierarchy levels:

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

The proxy address is the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy.

The remote address is the IP or MAC address of the remote router. The **remote** statement provides ARP capability from the TCC switching router to the Ethernet neighbor. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

Ethernet TCC is supported on interfaces that carry IPv4 traffic only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC and extended VLAN CCC are not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.

Example: Configuring an Ethernet TCC or Extended VLAN TCC

Configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. Ethernet TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. (See the topology in Figure 19 on page 234.)

The Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit accepts packets carrying standard TPID values.

Copyright © 2013, Juniper Networks, Inc.

If traffic flows from Router A to Router C, the Junos OS strips all PPP encapsulation data from incoming packets and adds Ethernet encapsulation data before forwarding the packets. If traffic flows from Router C to Router A, the Junos OS strips all Ethernet encapsulation data from incoming packets and adds PPP encapsulation data before forwarding the packets.

Figure 19: Topology of Layer 2.5 Translational Cross-Connect

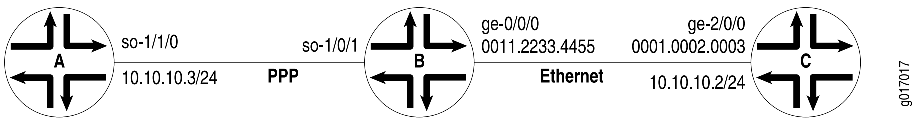

On Router B

```
interfaces
```

Configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. Extended VLAN TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. (See the topology in Figure 19 on page 234.)

Configuring an Extended VLAN TCC

The Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit is Ethernet with VLAN tagging enabled.

On Router B

```
interfaces
  ge-0/0/0 {
    vlan-tagging;
  encapsulation extended-vlan-tcc;
    unit 0 {
      vlan-id 1;
      family tcc {
        proxy {
          inet-address 10.10.10.3;
        }
        remote {
          inet-address 10.10.10.2;
        }
      }
    }
```

Related Documentation

- *encapsulation*
- *remote*
- *proxy*
- TCC and Layer 2.5 Switching Overview on page 231
- Configuring VLAN TCC Encapsulation on page 231
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 9

Configuring Static ARP Table Entries

Static ARP Table Entries Overview

For Fast Ethernet, Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure static ARP table entries, defining mappings between IP and MAC addresses.

Related Documentation

- Configuring Static ARP Table Entries on page 235

- *Ethernet Interfaces*

Configuring Static ARP Table Entries

To configure static ARP table entries, include the **arp** statement:

arp *ip-address* (mac | multicast-mac) *mac-address* <publish>;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, include the **multicast-mac** statement.

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

For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

Copyright © 2013, Juniper Networks, Inc.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

NOTE: The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

Example: Configuring Static ARP Table Entries

Configure two static ARP table entries on the router or switch's management interface:

```
[edit interfaces]
fxp0 {
    unit 0 {
        family inet {
            address 10.10.0.11/24 {
                arp 10.10.0.99 mac 0001.0002.0003;
                arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
            }
        }
    }
}
```

Related Documentation

- Management Ethernet Interface Overview on page 423
- *EX Series Switches Interfaces Overview*
- *Applying Policers*
- *Configuring an Unnumbered Interface*
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 10

Configuring Restricted and Unrestricted Proxy ARP

Restricted and Unrestricted Proxy ARP Overview

By default, the Junos OS responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet Interfaces, you can configure the router or switches to proxy-reply to the ARP requests using the restricted or unrestricted proxy ARP configuration.

You might want to configure restricted or unrestricted proxy ARP for routers that act as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.

NOTE: **From Junos OS Release 10.0 onward, Junos OS does not respond to proxy ARP requests with the default route 0.0.0.0. This behavior is in compliance with RFC 1027.**

Restricted Proxy ARP

Restricted proxy ARP enables the router or switch to respond to the ARP requests in which the physical networks of the source and target are not the same and the router or switch has an active route to the target address in the ARP request. The router does not reply if the target address is on the same subnet and the same interface as the ARP requestor.

Unrestricted Proxy ARP

Unrestricted proxy ARP enables the router or switch to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.

Copyright © 2013, Juniper Networks, Inc.

WARNING: If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2 reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.

NOTE: While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

Topology Considerations for Unrestricted Proxy ARP

In most situations, you should not configure the router or switch to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used. Figure 20 on page 238 and Figure 21 on page 239 show examples of situations in which you might want to configure unrestricted proxy ARP.

In Figure 20 on page 238, the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In Figure 21 on page 239, the Broadband Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

Figure 20: Edge Device Case for Unrestricted Proxy ARP

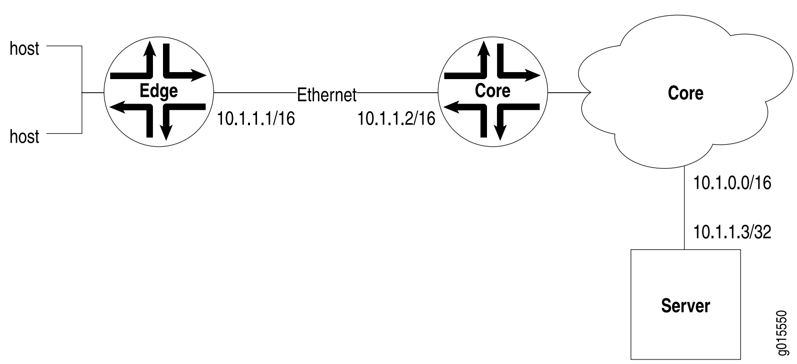

 Copyright © 2013, Juniper Networks, Inc.

Figure 21: Core Device Case for Unrestricted Proxy ARP

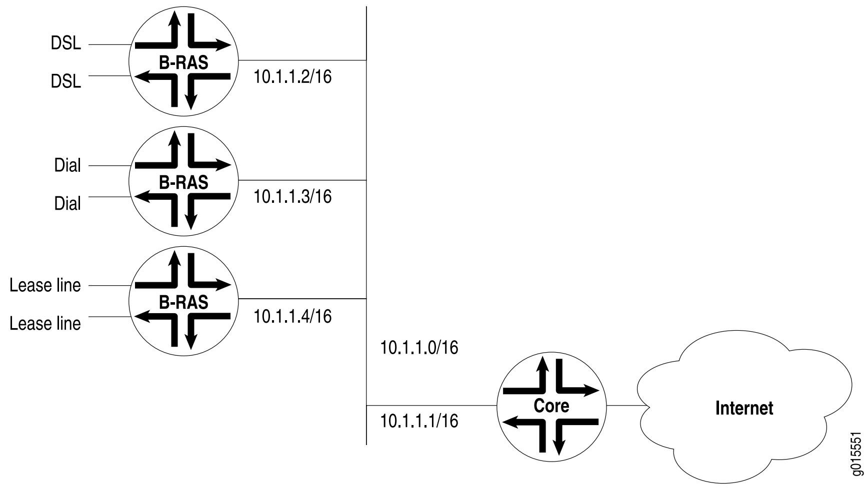

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 239
- *Ethernet Interfaces*

Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

```
proxy-arp (restricted |unrestricted);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.

Copyright © 2013, Juniper Networks, Inc.

NOTE: When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See "Configuring Gratuitous ARP" on page 48 for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- *proxy-arp*
- Restricted and Unrestricted Proxy ARP Overview on page 237
- Configuring Gratuitous ARP on page 48
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 11

Configuring MAC Address Validation on Static Ethernet Interfaces

MAC Address Validation on Static Ethernet Interfaces Overview

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

MAC address validation is supported on AE, Fast Ethernet, Gigabit Ethernet, and 10–Gigabit Ethernet interfaces (with or without VLAN tagging) on MX Series routers only.

There are two types of MAC address validation that you can configure:

- Loose—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

  Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not support the MAC address of the tuple

  Continues to forward packets when the source address of the incoming packet does not match any of the trusted IP addresses.
- Strict—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

  Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Related Documentation

- Configuring MAC Address Validation on Static Ethernet Interfaces on page 242
- Disabling MAC Address Learning of Neighbors Through ARP or Neighbor Discovery for IPv4 and IPv6 Neighbors on page 242
- *Ethernet Interfaces*

Configuring MAC Address Validation on Static Ethernet Interfaces

To configure MAC address validation on static Ethernet interfaces, include the **mac-validate (loose | strict)** statement in the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy:

```
[edit interfaces interface-name unit logical-unit-number family family]
mac-validate (loose | strict);
```

Example of Strict MAC Validation on a Static Ethernet Interface

This example shows strict MAC address validation on a static Ethernet interface without VLAN tagging.

```
[edit interfaces]
ge-2/1/9 {
  unit 0 {
    proxy-arp;
    family inet {
      mac-validate strict;
      address 88.22.100.1/24 {
        arp 88.22.100.3 mac 00:00:58:16:64:03;
      }
    }
  }
}
```

Related Documentation

- *family*
- *mac-validate*
- *Ethernet Interfaces*

Disabling MAC Address Learning of Neighbors Through ARP or Neighbor Discovery for IPv4 and IPv6 Neighbors

The Junos OS provides the **no-neighbor-learn** configuration statement at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** and **[edit interfaces *interface-name* unit *interface-unit-number* family inet6]** hierarchy levels.

To disable ARP address learning by not sending arp-requests and not learning from ARP replies for IPv4 neigbors, include the **no-neighbor-learn** statement at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet**] hierarchy level:

```
[edit interfaces interface-name unit interface-unit-number family inet]
no-neighbor-learn;
```

To disable neighbor discovery for IPv6 neighbors, include the **no-neighbor-learn** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level:

```
[edit interfaces interface-name unit interface-unit-number family inet6]
no-neighbor-learn;
```

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Configuring the Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 12

Enabling Passive Monitoring on Ethernet Interfaces

Passive Monitoring on Ethernet Interfaces Overview

The Monitoring Services I and Monitoring Services II PICs are designed to enable IP services. You can monitor IPv4 traffic if you have a Monitoring Services PIC installed in the router with the following PICs:

- 10-port Gigabit Ethernet PIC with SFPs
- 4-port Gigabit Ethernet PIC with SFPs
- 2-port Gigabit Ethernet PIC with SFPs
- 1-port 10-Gigabit Ethernet PIC

NOTE: The PICs in the preceding list support only IPv4.

NOTE: I2.0 based M120 routers and I3.0 based M320 routers with the PICs in the preceding list support passive monitoring starting with Junos OS Release 9.5. Other M Series and T Series routers with the PICs listed above started supporting passive monitoring before Junos OS Release 7.3. Support for 1-port 10-Gigabit Ethernet PIC with XENPAK on I2.0-based M120 routers and I3.0-based M320 routers was added in Junos OS Release 9.5.

- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (T640, T1600, and T4000 Core Routers) (supported on both WAN-PHY and LAN-PHY modes for both IPv4 and IPv6 addresses)

The following interfaces support passive monitoring on the I3.0-based MX 240, MX 480, and MX 960 routers, starting with Junos OS Release 8.5:

- Type 2 MX FPCs
- Type 3 MX FPCs
- Gigabit Ethernet Enhanced DPC with SFP (DPCE-R-40GE-SFP)
- 4-port 10-Gigabit Ethernet Enhanced DPCs with XFP (DPCE-R-4XGE-XFP)

The following interfaces support passive monitoring on the Trio-based MX 240, MX 480, and MX 960 routers:

- 10-Gigabit Ethernet MPC with SFP+
- 30-Gigabit Ethernet MPC
- 60-Gigabit Ethernet MPC

Passive monitoring is also supported on MX 80 routers with 10-Gigabit Ethernet MPC with SFP+ and 30-Gigabit Ethernet MPC interfaces.

Interfaces configured on the following FPCs and PIC support IPv6 passive monitoring on the T640, T1600, and T4000 routers:

- Enhanced Scaling FPC2
- Enhanced Scaling FPC3
- Enhanced Scaling FPC4
- Enhanced Scaling FPC4.1
- Enhanced II FPC1 (T640 and T1600 routers)
- Enhanced II FPC2 (T640 and T1600 routers)
- Enhanced II FPC3 (T640 and T1600 routers)
- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (supported on both WAN-PHY and LAN-PHY modes for both IPv4 and IPv6 addresses)
- Gigabit Ethernet PIC with SFP
- 10-Gigabit Ethernet PIC with XENPAK (T640 and T1600 routers)
- SONET/SDH OC192/STM64 PICs with XFP (T1600 and T4000 routers)
- SONET/SDH OC48c/STM16 PIC with SFP
- SONET/SDH OC12/STM4 (Multi-Rate) PIC with SFP (T1600 router)
- Type 1 SONET/SDH OC3/STM1 (Multi-Rate) PIC with SFP (T1600 router)

NOTE: **Unlike IPv4 passive monitoring, IPv6 passive monitoring is not supported on Monitoring Services PICs. You must configure port mirroring to forward the packets from the passive monitored ports to other interfaces.**

Related Documentation

- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Enabling Passive Monitoring on Ethernet Interfaces

On Ethernet interfaces, enable packet flow monitoring by including the **passive-monitor-mode** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
passive-monitor-mode;
```

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you include the **passive-monitor-mode** statement in the configuration:

- Gigabit and Fast Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.
- Ethernet interfaces do not support the **stacked-vlan-tagging** statement for both IPv4 and IPv6 packets in passive monitor mode.

For IPv4 monitoring services interfaces, enable packet flow monitoring by including the **family** statement at the **[edit interfaces mo-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level, specifying the **inet** option:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number]
family inet;
```

For conformity with the cflowd record structure, you must include the **receive-options-packets** and **receive-ttl-exceeded** statements at the **[edit interfaces mo-*fpc/pic/port* unit *logical-unit-number* family inet]** hierarchy level:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]
receive-options-packets;
receive-ttl-exceeded;
```

IPv6 passive monitoring is not supported on monitoring services PICs. A user must configure port mirroring to forward the packets from the passive monitored ports to other interfaces.

For information on FPCs and PICs that support IPv6 passive monitoring on the T640, T1600, and T4000 routers, see "Passive Monitoring on Ethernet Interfaces Overview" on page 245. Interfaces configured on these FPCs and PICs support IPv6 passive monitoring.

To configure port mirroring, include the **port-mirroring** statement at the **[edit forwarding-options]** hierarchy level.

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see *Configuring Multiservice Physical Interface Properties* and the *Junos OS Services Interfaces Library for Routing Devices*.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Passive Monitoring on Ethernet Interfaces Overview on page 245
- *Configuring Multiservice Physical Interface Properties*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 13

Configuring IEEE 802.1ag OAM Connectivity-Fault Management

IEEE 802.1ag OAM Connectivity Fault Management Overview

Ethernet interfaces on M7i and M10i routers with the Enhanced CFEB (CFEB-E) and on M120, M320, MX Series, T Series, and PTX Series routers support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). The goal of CFM is to monitor an Ethernet network that may comprise one or more service instances. Junos OS supports IEEE 802.1ag connectivity fault management.

In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces. On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet

Copyright © 2013, Juniper Networks, Inc.

member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces.
CFM does not support Multichassis Link Aggregation (MC-LAG). Do not configure the **mc-ae** statement when you configure CFM.

On T Series routers, CFM is not supported on interfaces configured with CCC encapsulation. If you configure CFM, the system displays the following message: "**MEPs cannot be configured on ccc interface on this platform**".

Network entities such as operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management, perform end-to-end monitoring, and still avoid security breaches.

NOTE: As a requirement for Ethernet OAM 802.1ag to work, distributed periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can only disable PPM on the Packet Forwarding Engine. To disable PPM on the PFE, include the **ppm no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

IEEE 802.1ag OAM supports graceful Routing Engine switchover (GRES). IEEE 802.1ag OAM is supported on untagged, single tagged, and stacked VLAN interfaces.

- Connectivity Fault Management Key Elements on page 250

Connectivity Fault Management Key Elements

Figure 22 on page 250 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 22: Relationship Among MEPs, MIPs, and Maintenance Domain Levels

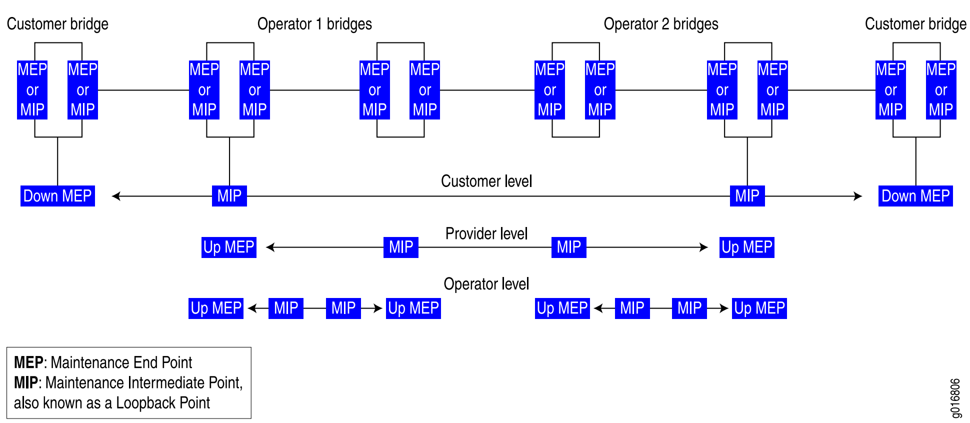

 Copyright © 2013, Juniper Networks, Inc.

NOTE: Maintenance intermediate points (MIP) are not supported on the ACX Series routers.

A maintenance association is a set of MEPs configured with the same maintenance association identifier and maintenance domain level. Figure 23 on page 251 shows the hierarchical relationships between the Ethernet bridge, maintenance domains, maintenance associations, and MEPs.

Figure 23: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs

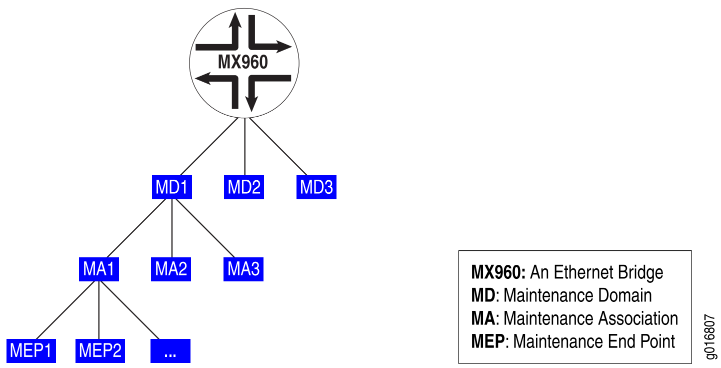

Related Documentation

- *connectivity-fault-management*
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Creating the Maintenance Domain

To enable CFM on an Ethernet interface, maintenance domains, maintenance associations, and MEPs must be created and configured.

To create a maintenance domain, include the **maintenance-domain *domain-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Give the maintenance domain a name. Names can be in one of several formats:

- Configuring the Maintenance Domain Name Format on page 252
- Configuring the Maintenance Domain Level on page 252

Configuring the Maintenance Domain Name Format

You can specify the maintenance domain name format as one of the following:

- A plain ASCII character string.
- A domain name service (DNS) format, a MAC address plus a two-octet identifier in the range from 0 through 65,535, or none.
- A MAC address plus a two-octet identifier in the range from 0 through 65,535.
- Or none.

If none is specified, the maintenance domain name is not used.

The default name format is an ASCII character string.

To configure the maintenance domain name format, include the **name-format (character-string | none | dns | mac+2octet)** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

NOTE: **If you configure the maintenance domain name length greater than 45 octet, then the following error message is displayed::**

error: configuration check-out failed

Configuring the Maintenance Domain Level

The maintenance domain level is a mandatory parameter that indicates the nesting relationship between various maintenance domains. The level is embedded in each of the CFM frames. CFM messages within a given level are processed by MEPs at that same level. For example, the operator domain can be level 0, the provider domain can be level 3, and the customer domain can be level 7.

To configure the maintenance domain level, include the **level *number*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *connectivity-fault-management*
- *maintenance-domain*

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296

- *Ethernet Interfaces*

Configuring Maintenance Intermediate Points

MX Series routers support maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag CFM protocol at a bridge-domain level. This enables you to define a maintenance domain for each default level. The MIPs names are created as **default-*level-number*** at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level. Use the **bridge-domain**, **instance**, **virtual-switch**, and **mip-half-function** MIP options to specify the MIP configuration.

NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations be the same.

To display MIP configurations, use the **show oam ethernet connectivity-fault-management mip (bridge-domain | instance-name | interface-name)** command.

The following sections describe MIP configuration:

- Configuring MIP for Bridge Domains of a Virtual Switch on page 254
- Configuring the Maintenance Domain Bridge Domain on page 254

- Configuring the Maintenance Domain Instance on page 254
- Configuring the Maintenance Domain MIP Half Function on page 254

Configuring MIP for Bridge Domains of a Virtual Switch

The default maintenance domain configuration allows MIP configuration for bridge domains for a default virtual switch or a user-defined virtual switch. You can use the **virtual-switch** and **bridge-domain** statements to specify which MIPs to enable for a user-defined virtual switch.

A bridge domain must be specified by name only if it is configured by including the **vlan-id *id*** statement under the **virtual-switch** statement.

If a bridge domain is configured with a range of VLAN IDs, then the VLAN IDs must be explicitly listed after the bridge domain name.

To configure a bridge domain under a user-defined virtual switch, include the **virtual-switch *name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* default-*x*]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  domain-name default-x]
virtual-switch name {
  bridge-domain {
    name-1;
    name-2 {
      vlan-id [vlan-ids ];
    }
  }
}
```

Configuring the Maintenance Domain Bridge Domain

The VLAN corresponds to the bridge domain.

To configure the bridge domain for the default virtual switch, include the **bridge-domain** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name*]** hierarchy level.

Configuring the Maintenance Domain Instance

To configure the maintenance domain instance for a VPLS routing instance, include the **instance** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level.

Configuring the Maintenance Domain MIP Half Function

MIP Half Function (MHF) divides MIP functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loopback and linktrace messages to help isolate faults.

Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the *MIP half function* value for

 Copyright © 2013, Juniper Networks, Inc.

all maintenance domains and maintenance associations be the same. To configure the MIP half function, include the **mip-half-function** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level.

Related Documentation

- *bridge-domain*
- *connectivity-fault-management*
- *instance*
- **mip-half-function on page 607**
- *virtual-switch*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

Creating a Maintenance Association

To create a maintenance association, include the **maintenance-association *ma-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

Maintenance association names can be in one of the following formats:

- As a plain ASCII character string
- As the VLAN identifier of the VLAN you primarily associate with the maintenance association
- As a two-octet identifier in the range from 0 through 65,535
- As a name in the format specified by RFC 2685

The default short name format is an ASCII character string.

Copyright © 2013, Juniper Networks, Inc.

To configure the maintenance association short name format, include the **short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id)** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level.

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

Continuity Check Protocol

The continuity check protocol is used for fault detection by a MEP within a maintenance association. The MEP periodically sends continuity check multicast messages. The receiving MEPs use the continuity check messages to build a MEP database of all MEPs in the maintenance association.

The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32.

- Configuring the Continuity Check on page 256
- Configuring the Continuity Check Hold Interval on page 257
- Configuring the Continuity Check Interval on page 257
- Configuring the Continuity Check Loss Threshold on page 258
- Continuity Measurement on page 258

Configuring the Continuity Check

You can configure the following continuity check protocol parameters:

- **hold-interval** ***minutes***

 Copyright © 2013, Juniper Networks, Inc.

- **interval** *time*
- **loss-threshold** *number*

To enable the continuity check protocol, include the **continuity-check** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name*****]** hierarchy level.

Configuring the Continuity Check Hold Interval

You can specify the continuity check hold interval. The hold interval is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

The hold interval logic runs a polling timer per CFM session level (not per remote MEP level) where the polling timer duration is equal to the configured hold time. When the polling timer expires, it deletes all the auto discovered remote MEP entries which have been in the failed state for a time period equal to or greater than the configured hold time. If the remote MEP completes the hold time duration in the failed state, then flushing will not occur until the next polling timer expires. Hence remote MEP flushing may not happen exactly at the configured hold time.

To configure the hold interval, include the **hold-interval** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name*** **continuity-check]** hierarchy level.

NOTE: Hold timer based flushing is applicable only for auto discovered remote MEPs and not for statically configured remote MEPs.

Configuring the Continuity Check Interval

You can specify the continuity check message (CCM) interval. The interval is the time between the transmission of CCMs. You can specify 10 minutes (**10m**), 1 minute (**1m**), 10 seconds (**10s**), 1 second (**1s**), 100 milliseconds (**100ms**), or 10 milliseconds (**10ms**). The default value is 1 minute.

NOTE: For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine (PFE) by default. You can only disable PPM on the PFE. To disable PPM on the PFE, use the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

Continuity check interval of 10 milliseconds is not supported for CFM sessions over a Label-Switched interface (LSI).

To configure the interval, include the **interval** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain** ***domain-name*** **maintenance-association** ***ma-name*** **continuity-check]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Configuring the Continuity Check Loss Threshold

You can specify the number of continuity check messages that can be lost before marking the MEP as down. The default value is three (PDUs).

To configure the loss threshold, include the **loss-threshold** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* continuity-check]** hierarchy level.

Continuity Measurement

Continuity measurement is provided by an existing continuity check protocol. The continuity for every remote MEP is measured as the percentage of time that remote MEP was operationally up over the total administratively enabled time. Here, the operational uptime is the total time during which the CCM adjacency is active for a particular remote MEP and the administrative enabled time is the total time during which the local MEP is active. You can also restart the continuity measurement by clearing the currently measured operational uptime and the administrative enabled time.

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- Managing Continuity Measurement Statistics on page 358
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Configuring a Maintenance Endpoint

To configure the maintenance endpoint, include the **mep** ***mep-id*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level.

Enabling Maintenance Endpoint Automatic Discovery

You can enable the MEP to accept continuity check messages from all remote MEPs of the same maintenance association.

To configure automatic discovery, include the **auto-discovery** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Configuring the Maintenance Endpoint Direction

You can specify the direction in which CFM packets are transmitted for the MEP.

Direction up continuity check messages (CCMs) are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on this MEP.

Direction down CCMs are transmitted only out of the interface configured on this MEP.

NOTE: **Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.**

To configure the MEP direction, include the **direction** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 3D Universal Edge Routers, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all our routers and switches, you must continue to configure the **no-control-word** statement at the **[edit routing-instances routing-instance-name protocols l2vpn]** or **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

Configuring the Maintenance Endpoint Interface

You must specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface.

On MX Series routers, you can enable the MEP on a specific VLAN of a trunk interface.

To configure the interface, include the **interface *interface-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

MEP Interface Configuration

This example shows the MEP interface configuration statements:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  domain-name maintenance-association ma-name]
mep mep-id {
  direction (up | down);
  interface (ge | xe)-(fpc/pic/port | fpc/pic/port.domain | fpc/pic/port.domain vlan vlan-id);
  auto-discovery;
  priority number;
}
```

Configuring the Maintenance Endpoint Priority

You can specify the IEEE 802.1 priority bits that are used by continuity check and link trace messages.

To configure the priority, include the **priority** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Configuring the Maintenance Endpoint Lowest Priority Defect

You can specify the lowest priority defect that is allowed to generate a fault alarm. This configuration determines whether to generate a fault alarm whenever it detects a defect. This configuration is done at the MEP level.

Copyright © 2013, Juniper Networks, Inc.

To configure the lowest priority defect, include the **lowest-priority-defect *options*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Table 16 on page 261 describes the available lowest priority defect options.

Table 16: Lowest Priority Defect Options

Option	Description
all-defects	Allows all defects.
err-xcon	Allows only erroneous CCM and cross-connect CCM defects.
mac-rem-err-xcon	Allows only MAC, not receiving CCM, erroneous CCM, and cross-connect defects.
no-defect	Allows no defect.
rem-err-xcon	Allows only not receiving CCM, erroneous CCM, and cross-connect CCM defects.
xcon	Allows only cross-connect CCM defects.

The following configuration example shows **mac-rem-err-xcon** as the lowest priority defect:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md6 {
        level 6;
        maintenance-association ma6 {
          mep 200 {
            interface ge-5/0/0.0;
            direction down;
            lowest-priority-defect mac-rem-err-xcon;
          }
        }
      }
    }
  }
}
```

Configuring a Remote Maintenance Endpoint

You can configure a remote MEP from which CCM messages are expected. If autodiscovery is not enabled, the remote MEP must be configured under the **mep** statement. If the remote MEP is not configured under the **mep** statement, the CCMs from the remote MEP are treated as errors.

To configure the remote MEP, include the **remote-mep** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]** hierarchy level.

Copyright © 2013, Juniper Networks, Inc.

Configuring a Remote Maintenance Endpoint Action Profile

You can specify the name of the action profile to use for the remote MEP.

To configure the action profile, include the **action-profile *profile-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id*]** hierarchy level. The profile must already be defined at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Configuring Maintenance Endpoint Service Protection

You can enable service protection for a VPWS (Virtual Private Wire Service) over MPLS by specifying a working path or protect path on the MEP. Service protection provides end-to-end connection protection of the working path in the event of a failure.

To configure service protection, you must create two separate transport paths a working path and a protect path. You can specify the working path and protect path by creating two maintenance associations. To associate the maintenance association with a path, you must configure the MEP **interface** statement within the maintenance association and specify the path as working or protect.

To configure the MEP interface, include the **interface** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*** hierarchy level. On the **interface** statement, specify the path as (**working | protect**). The direction must also be configured as direction down for both sessions.

NOTE: **If the path is not specified, the session monitors the active path.**

Table 17 on page 262 describes the available service protection options.

Table 17: Service Protection Options

Option	Description
working	Specifies the working path.
protect	Specifies the protect path.

The following configuration example shows service protection is enabled for the VPWS service. The CCM session is configured for the working path and references the CCM session configured for the protect path in the **protect-maintenance-association** statement. The APS profile is configured and associated with the maintenance-association for the working path:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
```

 Copyright © 2013, Juniper Networks, Inc.

```
            maintenance-domain vpws-service-1 {
                name-format none;
                level 5;
                maintenance-association W {
                    short-name-format character-string;
                        protect-maintenance-association P {
                        aps-profile aps-profile-1;
                    }
                    continuity-check {
                        interval 1s;
                    }
                    mep 1 {
                        interface ge-1/3/5.0 working;
                        direction down;
                        auto-discovery;
                    }
                }
                maintenance-association P {
                    short-name-format character-string;
                    continuity-check {
                        interval 1s;
                    }
                    mep 1 {
                        interface ge-1/3/5.0 protect;
                        direction down;
                        auto-discovery;
                    }
                }
            }
        }
    }
}
```

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296

Copyright © 2013, Juniper Networks, Inc.

- *Ethernet Interfaces*

Configuring a Connectivity Fault Management Action Profile

You can configure an action profile and specify the action to be taken when any of the configured events occur. Alternatively, you can configure an action profile and specify default actions when connectivity to a remote maintenance association endpoint (MEP) fails.

To configure the action profile name, include the **action-profile** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Configuring the Action of a CFM Action Profile

You can configure the action to be taken when any of the configured events occur.

To configure the action profile's action, include the **action** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name*]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile bring-down {
      event {
        interface-status-tlv down;
      }
      action {
        interface-down;
      }
    }
  }
}
```

Configuring the Default Actions of a CFM Action Profile

You can configure the default actions to be taken when connectivity to a remote MEP fails.

To enable the **interface-down** as the default action for an action profile, include the **interface-down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* default-actions]** hierarchy level.

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
```

 Copyright © 2013, Juniper Networks, Inc.

```
                action-profile bring-down {
                    default-actions {
                        interface-down;
                    }
                }
                maintenance-domain md1 {
                    level 0;
                        maintenance-association ma1 {
                        continuity-check {
                            interval 100 ms;
                        }
                        mep 4001 {
                            interface ge-4/1/0;
                            direction down;
                            remote-mep 1 {
                                action-profile bring-down;
                            }
                        }
                    }
                }
            }
        }
    }
}
```

NOTE: Associating an **action-profile** with the action of **interface-down** on an up MEP CFM session running over a circuit cross-connect (CCC) interface (l2circuit/l2vpn) is not advisable and can result in a deadlock situation.

Configuring a CFM Action Profile Event

You can configure one or more events under the action profile, the occurrence of which triggers the corresponding action to be taken.

To configure the interface-status-tlv lower-layer-down event, include the **interface-status-tlv lower-layer-down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure the interface-status-tlv down event, include the **interface-status-tlv down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure the port-status-tlv blocked event, include the **port-status-tlv blocked** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure the adjacency-loss event, include the **adjacency-loss** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure an RDI event to bring an interface down on reception of an RDI bit from a MEP, include the **rdi** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile bring-down {
      event {
        adjacency-loss;
        interface-status-tlv (down | lower-layer-down);
        port-status-tlv blocked;
        rdi;
      }
      action {
        interface-down;
      }
      clear-action {
        interface-down peer-interface ;
      }
    }
  }
}
```

NOTE: You cannot configure multiple actions at this time. Only one action can be configured. This limitation affects both the **action** and **clear-action** statements.

Related Documentation

- *event (CFM)*
- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Creating a Maintenance Association on page 255
- Configuring a Maintenance Endpoint on page 259
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Configuring Linktrace Protocol in CFM

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between an MEP and an MIP under the same maintenance domain. The operation of IEEE 802.1ag linktrace request and response messages is similar to the operation of Layer 3 **traceroute** commands. For more information about the **traceroute** command, see the *Junos OS Administration Library for Routing Devices*.

Configuring the Linktrace Path Age Timer

If no response to a **linktrace** request is received, the request and response entries are deleted after the age timer expires. To configure the linktrace age timer, use the **linktrace** statement with the **age *time*** option at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level. The age is configured in minutes or seconds.

Configuring the Linktrace Database Size

Configure the number of linktrace reply entries to be stored per linktrace request. To configure the linktrace database size, use the **linktrace** statement with the **path-database-size *path-database-size*** option at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

Display the linktrace database using the **show oam ethernet connectivity-fault-management path-database** command.

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296

Copyright © 2013, Juniper Networks, Inc.

- *Ethernet Interfaces*

Configuring Ethernet Local Management Interface

Ethernet Local Management Interface Overview

MX Series routers with Gigabit Ethernet (**ge**), 10-Gigabit Ethernet (**xe**), or Aggregated Ethernet (**ae**) interfaces support the Ethernet Local Management Interface (E-LMI). The E-LMI specification is available at the Metro Ethernet Forum. E-LMI procedures and protocols are used for enabling automatic configuration of the customer edge (CE) to support Metro Ethernet services. The E-LMI protocol also provides user-to-network interface (UNI) and Ethernet virtual connection (EVC) status information to the CE. The UNI and EVC information enables automatic configuration of CE operation based on the Metro Ethernet configuration.

The E-LMI protocol operates between the CE device and the provider edge (PE) device. It runs only on the PE-CE link and notifies the CE of connectivity status and configuration parameters of Ethernet services available on the CE port. The scope of the E-LMI protocol is shown in Figure 24 on page 268.

Figure 24: Scope of the E-LMI Protocol

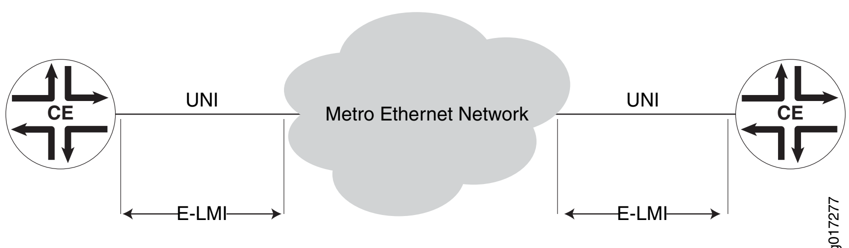

The E-LMI implementation on MX Series routers includes only the PE side of the E-LMI protocol.

E-LMI interoperates with an OAM protocol, such as Connectivity Fault Management (CFM), that runs within the provider network to collect OAM status. CFM runs at the provider maintenance level (UNI-N to UNI-N with up MEPs at the UNI). E-LMI relies on the CFM for end-to-end status of EVCs across CFM domains (SVLAN domain or VPLS).

The E-LMI protocol relays the following information:

- Notification to the CE of the addition/deletion of an EVC (active, not active, or partially active)
- Notification to the CE of the availability state of a configured EVC
- Communication of UNI and EVC attributes to the CE:

 Copyright © 2013, Juniper Networks, Inc.

- UNI attributes:
  - UNI identifier (a user-configured name for UNI)
  - CE-VLAN ID/EVC map type (all-to-one bundling, service multiplexing with bundling, or no bundling)
  - Bandwidth profile is not supported (including the following features):
    - CM (coupling mode)
    - CF (color flag)
    - CIR (committed Information rate)
    - CBR (committed burst size)
    - EIR (excess information rate)
    - EBS (excess burst size)
- EVC attributes:
  - EVC reference ID
  - EVC status type (active, not active, or partially active)
  - EVC type (point-to-point or multipoint-to-multipoint)
  - EVC ID (a user-configured name for EVC)
  - Bandwidth profile (not supported)
- CE-VLAN ID/EVC map

E-LMI on MX Series routers supports the following EVC types:

- Q-in-Q SVLAN (point-to-point or multipoint-to-multipoint)—Requires an end-to-end CFM session between UNI-Ns to monitor the EVS status.
- VPLS (BGP or LDP) (point-to-point or multipoint-to-multipoint)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.
- L2 circuit/L2VPN (point-to-point)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.

NOTE: **l2-circuit and l2vpn are not supported.**

Configuring the Ethernet Local Management Interface

To configure E-LMI, perform the following steps:

- Configuring an OAM Protocol (CFM) on page 270
- Assigning the OAM Protocol to an EVC on page 270
- Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC on page 270

Copyright © 2013, Juniper Networks, Inc.

Configuring an OAM Protocol (CFM)

For information on configuring the OAM protocol (CFM), see "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 249.

Assigning the OAM Protocol to an EVC

To configure an EVC, you must specify a name for the EVC using the **evcs*evc-id*** statement at the **[edit protocols oam ethernet]** hierarchy level. You can set the EVC protocol for monitoring EVC statistics to **cfm** or **vpls** using the **evc-protocol** statement and its options at the **[edit protocols oam ethernet evcs]** hierarchy level.

You can set the number of remote UNIs in the EVC using the **remote-uni-count *number*** statement at the **[edit protocols oam ethernet evcs evcs-protocol]** hierarchy level. The **remote-uni-count** defaults to 1. Configuring a value greater than 1 makes the EVC multipoint-to-multipoint. If you enter a value greater than the actual number of endpoints, the EVC status will display as partially active even if all endpoints are up. If you enter a **remote-uni-count** less than the actual number of endpoints, the status will display as active, even if all endpoints are not up.

You can configure an EVC by including the **evcs** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
evcs evc-id {
  evc-protocol (cfm (management-domain name management-association name ) | vpls
    (routing-instance name)) {
    remote-uni-count <number>; # Optional, defaults to 1
    multipoint-to-multipoint;
      # Optional, defaults to point-to-point if remote-uni-count is 1
  }
}
```

Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC

To configure E-LMI, include the **lmi** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
lmi {
  polling-verification-timer value;
    # Polling verification timer (T392), defaults to 15 seconds
  status-counter count; # Status counter (N393), defaults to 4
  interface name {
    evc evc-id {
      default-evc;
      vlan-list [ vlan-ids ];
    }
    evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
    polling-verification-time value; # Optional, defaults to global value
    status-counter count; # Optional, defaults to global value
    uni-id value; # Optional, defaults to interface-name
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

You can set the status counter to count consecutive errors using the **status-counter** ***count*** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The status counter is used to determine if E-LMI is operational or not. The default value is 4.

You can set the **polling-verification-timer** ***value*** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The default value is 15 seconds.

You can enable an interface and set its options for use with E-LMI using the **interface** ***name*** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. Only **ge**, **xe**, and **ae** interfaces are supported. You can use the interface **uni-id** option to specify a name for the UNI. If **uni-id** is not configured, it defaults to the name variable of **interface** ***name***.

You can specify the CE-VLAN ID/EVC map type using the **evc-map-type type** interface option. The options are **all-to-one-bundling**, **bundling**, or **service-multiplexing**. Service multiplexing is with no bundling. The default type is **all-to-one-bundling**.

To specify the EVC that an interface uses, use the **evc** ***evc-id*** statement at the **[edit protocols oam ethernet lmi interface** ***name*****]** hierarchy level. You can specify an interface as the default EVC interface using the **default-evc** statement at the **[edit protocols oam ethernet lmi interface** ***name*** **evc** ***evc-id*****]** hierarchy level. All VIDs that are not mapped to any other EVCs are mapped to this EVC. Only one EVC can be configured as the default.

You can map a list of VLANs to an EVC using the **vlan-list** ***vlan-id-list*** statement at the **[edit protocols oam ethernet lmi interface** ***name*** **evc** ***evc-id*****]** hierarchy level.

Example E-LMI Configuration

Figure 25 on page 271 illustrates the E-LMI configuration for a point-to-point EVC (SVLAN) monitored by CFM. In this example, VLANs 1 through 2048 are mapped to **evc1** (SVLAN 100) and 2049 through 4096 are mapped to **evc2** (SVLAN 200). Two CFM sessions are created to monitor these EVCs.

Figure 25: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM

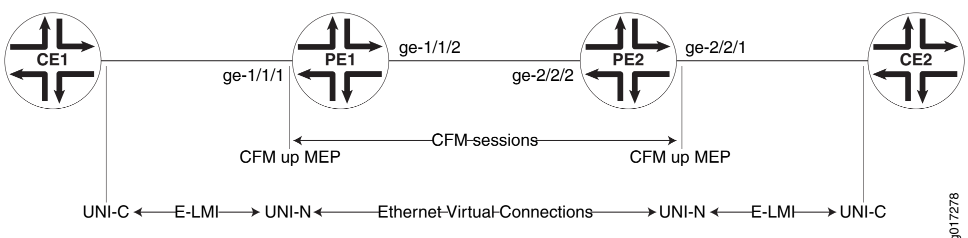

Configuring PE1

```
[edit]
interfaces {
   ge-1/1/1 {
      unit 0 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 1-2048;
         }
```

```
            }
            unit 1 {
                family bridge {
                    interface-mode trunk;
                    vlan-id-list 2049-4096;
                }
            }
        }
        ge-1/1/2 {
            unit 0 {
                vlan-id 100;
                family bridge {
                    interface-mode trunk;
                    inner-vlan-id-list 1-2048;
                }
            }
            unit 1 {
                vlan-id 200;
                family bridge {
                    interface-mode trunk;
                    inner-vlan-id-list 2049-4096;
                }
            }
        }
    }
    protocols {
        oam {
            ethernet {
                connectivity-fault-management {
                    maintenance-domain md {
                        level 0;
                        maintenance-association 1 {
                            name-format vlan;
                            mep 1 {
                                direction up;
                                interface ge-1/1/1.0 vlan 1;
                            }
                        }
                        maintenance-association 2049 {
                            name-format vlan;
                            mep 1 {
                                direction up;
                                interface ge-1/1/1.1 vlan 2049;
                            }
                        }
                    }
                }
                evcs {
                    evc1 {
                        evc-protocol cfm management-domain md management-association 1;
                        remote-uni-count 1;
                    }
                    evc2 {
                        evc-protocol cfm management-domain md management-association 2049;
                        remote-uni-count 1;
                    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
            }
            lmi {
               interface ge-1/1/1 {
                  evc evc1 {
                     vlan-list 1-2048;
                  }
                  evc evc2 {
                     vlan-list 2049-4096;
                  }
                  evc-map-type bundling;
                  uni-id uni-ce1;
               }
            }
         }
      }
   }
```

Configuring PE2

```
[edit]
interfaces {
   ge-2/2/1 {
      unit 0 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 1-2048;
         }
      }
      unit 1 {
         family bridge {
            interface-mode trunk;
            vlan-id-list 2049-4096;
         }
      }
   }
   ge-2/2/2 {
      unit 0 {
         vlan-id 100;
         family bridge {
            interface-mode trunk;
            inner-vlan-id-list 1-2048;
         }
      }
      unit 1 {
         vlan-id 200;
         family bridge {
            interface-mode trunk;
            inner-vlan-id-list 2049-4095;
         }
      }
   }
}
protocols {
   oam {
      ethernet {
         connectivity-fault-management {
```

Copyright © 2013, Juniper Networks, Inc.

```
                maintenance-domain md {
                    level 0;
                    maintenance-association 1 {
                        name-format vlan;
                        mep 1 {
                            direction up;
                            interface ge-2/2/1.0 vlan 1;
                        }
                    }
                    maintenance-association 2049 {
                        name-format vlan;
                        mep 1 {
                            direction up;
                            interface ge-2/2/1.1 vlan 2049;
                        }
                    }
                }
            }
            evcs {
                evc1 {
                    evc-protocol cfm management-domain md management-association 1;
                    remote-uni-count 1;
                }
                evc2 {
                    evc-protocol cfm management-domain md management-association 2049;
                    uni-count 2;
                }
            }
            lmi {
                interface ge-2/2/1 {
                    evc evc1 {
                        vlan-list 1-2048;
                    }
                    evc evc2 {
                        vlan-list 2049-4095;
                    }
                    evc-map-type bundling;
                    uni-id uni-ce2;
                }
            }
        }
    }
}
```

Configuring Two UNIs Sharing the Same EVC

```
[edit protocols]
oam {
    ethernet {
        connectivity-fault-management { ...}
        evcs {
            evc1 {
                evc-protocol cfm management-domain md management-association 1;
                remote-uni-count 1;
            }
        }
```


Copyright © 2013, Juniper Networks, Inc.

```
        lmi {
          interface ge-2/2/1 {
            evc evc1 {
              vlan-list 0-4095;
            }
            evc-map-type all-to-one-bundling;
            uni-id uni-ce1;
          }
          interface ge-2/3/1 {
            evc evc1 {
              vlan-list 0-4095;
            }
            evc-map-type all-to-one-bundling;
            uni-id uni-ce2;
          }
        }
      }
    }
```

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

Configuring Port Status TLV and Interface Status TLV

Copyright © 2013, Juniper Networks, Inc.

TLVs Overview

Type, Length, and Value (TLVs) are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and/or optional information in a PDU. TLVs are not aligned to any particular word or octet boundary. TLVs follow each other with no padding between them.

Table 18 on page 276 shows the TLV format and indicates if it is required or optional.

Table 18: Format of TLVs

Parameter	Octet (sequence)	Description
Type	1	Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.
Length	2–3	Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field.
Value	4	Length specified by the Length field. Optional. Not present if the Type field is 0 or if the Length field is 0.

Various TLVs for CFM PDUs

Table 19 on page 276 shows a set of TLVs defined by IEEE 802.1ag for various CFM PDU types. Each TLV can be identified by the unique value assigned to its type field. Some type field values are reserved.

Table 19: Type Field Values for Various TLVs for CFM PDUs

TLV or Organization	Type Field
End TLV	0
Sender ID TLV	1
Port Status TLV	2
Data TLV	3
Interface Status TLV	4
Reply Ingress TLV	5
Reply Egress TLV	6
LTM Egress Identifier TLV	7
LTR Egress Identifier TLV	8
Reserved for IEEE 802.1	9 to 30

 Copyright © 2013, Juniper Networks, Inc.

Table 19: Type Field Values for Various TLVs for CFM PDUs *(continued)*

TLV or Organization	Type Field
Organization-Specific TLV	31
Defined by ITU-T Y.1731	32 to 63
Reserved for IEEE 802.1	64 to 255

Not every TLV is applicable for all types of CFM PDUs.

- TLVs applicable for continuity check message (CCM):
  - End TLV
  - Sender ID TLV
  - Port Status TLV
  - Interface Status TLV
  - Organization-Specific TLV
- TLVs applicable for loopback message (LBM):
  - End TLV
  - Sender ID TLV
  - Data TLV
  - Organization-Specific TLV
- TLVs applicable for loopback reply (LBR):
  - End TLV
  - Sender ID TLV
  - Data TLV
  - Organization-Specific TLV
- TLVs applicable for linktrace message (LTM):
  - End TLV
  - LTM Egress Identifier TLV
  - Sender ID TLV
  - Organization-Specific TLV
- TLVs applicable for linktrace reply (LTR):
  - End TLV
  - LTR Egress Identifier TLV
  - Reply Ingress TLV

- Reply Egress TLV
- Sender ID TLV
- Organization-Specific TLV

The following TLVs are currently supported in the applicable CFM PDUs:

- End TLV
- Reply Ingress TLV
- Reply Egress TLV
- LTR Egress Identifier TLV
- LTM Egress Identifier TLV
- Data TLV

Support for Additional Optional TLVs

The following additional optional TLVs are supported:

- Port Status TLV
- Interface Status TLV

MX Series routers support configuration of port status TLV and interface status TLV. Configuring the Port Status TLV allows the operator to control the transmission of the Port Status TLV in CFM PDUs.

NOTE: Although Port Status TLV configuration statements are visible in the CLI on M120 and M320 routers, Port Status TLV cannot be configured on these systems. Port Status TLV can be enabled on a MEP interface only if it is a bridge logical interface, which is not possible on these systems.

For configuration information, see the following sections:

- Port Status TLV on page 278
- Interface Status TLV on page 281

Port Status TLV

The Port Status TLV indicates the ability of the bridge port on which the transmitting MEP resides to pass ordinary data, regardless of the status of the MAC. The value of this TLV is driven by the MEP variable **enableRmepDefect**, as shown in Table 21 on page 279. The format of this TLV is shown in Table 20 on page 279.

Any change in the Port Status TLVs value triggers one extra transmission of that bridge ports MEP CCMs.

 Copyright © 2013, Juniper Networks, Inc.

Table 20: Port Status TLV Format

Parameter	Octet (Sequence)
Type = 2	1
Length	2–3
Value (See Table 21 on page 279)	4

Table 21: Port Status TLV Values

Mnemonic	Ordinary Data Passing Freely Through the Port	Value
psBlocked	No: **enableRmepDefect** = false	1
psUp	Yes: **enableRmepDefect** = true	2

The MEP variable **enableRmepDefect** is a boolean variable indicating whether frames on the service instance monitored by the maintenance associations if this MEP are enabled to pass through this bridge port by the Spanning Tree Protocol and VLAN topology management. It is set to TRUE if:

- The bridge port is set in a state where the traffic can pass through it.
- The bridge port is running multiple instances of the spanning tree.
- The MEP interface is not associated with a bridging domain.

Configuring Port Status TLV

Junos OS provides configuration support for the Port Status TLV, allowing you to control the transmission of this TLV in CCM PDUs. The Junos OS provides this configuration at the continuity-check level. By default, the CCM does not include the Port Status TLV. To configure the Port Status TLV, use the **port-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check]** hierarchy level.

NOTE: Port Status TLV configuration is not mandated by IEEE 802.1ag. The Junos OS provides it in order to give more flexibility to the operator; however it receives and processes CCMs with a Port Status TLV, regardless of this configuration.

An example of the configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
        level number;
```

```
            maintenance-association identifier {
              continuity-check {
                interval number,
                loss-threshold number;
                hold-interval number;
                port-status-tlv; # Sets Port Status TLV
              }
            }
          }
        }
      }
    }
  }
```

You cannot enable Port Status TLV transmission in the following two cases:

- If the MEP interface under the maintenance-association is not of type bridge.
- If the MEP is configured on a physical interface.

Displaying the Received Port Status TLV

The Junos OS saves the last received Port Status TLV from a remote MEP. If the received Port Status value does not correspond to one of the standard values listed in Table 21 on page 279, then the **show** command displays it as "unknown." You can display the last saved received Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none # RX PORT STATUS
    Interface status TLV: none
```

Displaying the Transmitted Port Status TLV

The Junos OS saves the last transmitted Port Status TLV from a local MEP. If the transmission of the Port Status TLV has not been enabled, then the **show** command displays "none." You can display the last saved transmitted Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain**

 Copyright © 2013, Juniper Networks, Inc.

identifier maintenance-association *identifier* local-mep *identifier* remote-mep *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up # TX PORT STATUS
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

Interface Status TLV

The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the IETF RFC 2863 IF-MIB. The format of this TLV is shown in Table 22 on page 281. The enumerated values are shown in Table 23 on page 281.

Table 22: Interface Status TLV Format

Parameter	Octet (Sequence)
Type = 4	1
Length	2–3
Value (See Table 23 on page 281)	4

Table 23: Interface Status TLV Values

Mnemonic	Interface Status	Value
isUp	up	1
isDown	down	2
isTesting	testing	3
isUnknown	unknown	4
isDormant	dormant	5
isNotPresent	notPresent	6

Copyright © 2013, Juniper Networks, Inc.

Table 23: Interface Status TLV Values *(continued)*

Mnemonic	Interface Status	Value
isLowerLayerDown	lowerLayerDown	7

NOTE: When the operational status of a logical interface changes from the down state (status value of 2) to the lower layer down state (status value of 7) and vice versa, the LinkDown SNMP trap is not generated. For example, if you configure an aggregated Ethernet interface bundle with a VLAN tag and add a physical interface that is in the operationally down state to the bundle, the operational status of the aggregated Ethernet logical interface bundle at that point is lower layer down (7). If you take the MIC associated with the interface offline, the LinkDown trap is not generated when the logical interface shifts from the lower layer down state to the down state.

Similarly, consider another sample scenario in which an physical interface is added to an aggregated Ethernet bundle that has VLAN tagging and the aggregated Ethernet logical interface is disabled. When the logical interface is disabled, the operational status of the logical interface changes to down. If you disable the physical interface that is part of the aggregated Ethernet bundle, the operational status of the aggregated Ethernet logical interface remains down. If you reenable the aggregated Ethernet logical interface, the operational status of it changes from down to lower layer down. The LinkDown SNMP trap is not generated at this point.

Configuring Interface Status TLV

The Junos OS provides configuration support for the Interface Status TLV, thereby allowing operators to control the transmission of this TLV in CCM PDUs through configuration at the continuity-check level.

NOTE: This configuration is not mandated by IEEE 802.1ag; rather it is provided to give more flexibility to the operator. The Junos OS receives and processes CCMs with the Interface Status TLV, regardless of this configuration.

The interface status TLV configuration is shown below:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
          level number;
          maintenance-association identifier {
            continuity-check {
              interval number;
```

 Copyright © 2013, Juniper Networks, Inc.

```
                    loss-threshold number;
                    hold-interval number;
                    interface-status-tlv; # Sets the interface status TLV
                }
              }
            }
          }
        }
      }
    }
```

NOTE: The Junos OS supports transmission of only three out of seven possible values for the Interface Status TLV. The supported values are 1, 2, and 7. However, the Junos OS is capable of receiving any value for the Interface Status TLV.

Displaying the Received Interface Status TLV

The Junos OS saves the last received Interface Status TLV from the remote MEP. If the received Interface Status value does not correspond to one of the standard values listed in Table 22 on page 281, then the **show** command displays "unknown."

You can display this last saved Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain *md5* maintenance-association *ma5* local-mep *2001* remote-mep *1001***

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none # displays the Interface Status TLV state
```

Displaying the Transmitted Interface Status TLV

The Junos OS saves the last transmitted Interface Status TLV from a local MEP. If the transmission of Interface Status TLV has not been enabled, then the **show** command displays "none."

You can display the last transmitted Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier***

Copyright © 2013, Juniper Networks, Inc.

maintenance-association ***identifier*** **local-mep** ***identifier*** **remote-mep** ***identifier*** command, as in the following example:

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***md5*** **maintenance-association** ***ma5*** **local-mep** ***2001*** **remote-mep** ***1001***

```
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 1001, State: ok
    MAC address: 00:19:e2:b0:74:00, Type: Learned
    Interface: ge-2/0/0.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

MAC Status Defects

The Junos OS provides MAC status defect information, indicating that one or more of the remote MEPs is reporting a failure in its Port Status TLV or Interface Status TLV. It indicates "yes" if either some remote MEP is reporting that its interface is not isUp (for example, at least one remote MEPs interface is unavailable), or if all remote MEPs are reporting a Port Status TLV that contains some value other than psUp (for example, all remote MEPs Bridge Ports are not forwarding data). There are two **show** commands you can use to view the MAC Status Defects indication.

Use the **mep-database** command to display MAC status defects:

user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***md6*** **maintenance-association** ***ma6***

```
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: up, Port status TLV: up
  Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM            : no
    Erroneous CCM received                  : no
    Cross-connect CCM received              : no
    RDI sent by some MEP                    : no
    Some remote MEP's MAC in error state    : yes # MAC Status Defects yes/no
  Statistics:
    CCMs sent                               : 1658
    CCMs received out of sequence           : 0
    LBMs sent                               : 0
    Valid in-order LBRs received            : 0
    Valid out-of-order LBRs received        : 0
    LBRs received with corrupted data       : 0
    LBRs sent                               : 0
    LTMs sent                               : 0
```

 Copyright © 2013, Juniper Networks, Inc.

```
    LTMs received                          : 0
    LTRs sent                              : 0
    LTRs received                          : 0
    Sequence number of next LTM request    : 0
    1DMs sent                              : 0
    Valid 1DMs received                    : 0
    Invalid 1DMs received                  : 0
    DMMs sent                              : 0
    DMRs sent                              : 0
    Valid DMRs received                    : 0
    Invalid DMRs received                  : 0
  Remote MEP count: 1
    Identifier     MAC address        State    Interface
      200      00:05:85:73:39:4a       ok    xe-5/0/0.0
```

Use the **interfaces** command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
    Interface status TLV: up, Port status TLV: up
    MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                : no
      Erroneous CCM received                      : no
      Cross-connect CCM received                  : no
      RDI sent by some MEP                        : no
      Some remote MEP's MAC in error state        : yes # MAC Status Defects
yes/no
    Statistics:
      CCMs sent                                   : 1328
      CCMs received out of sequence               : 0
      LBMs sent                                   : 0
      Valid in-order LBRs received                : 0
      Valid out-of-order LBRs received            : 0
      LBRs received with corrupted data           : 0
      LBRs sent                                   : 0
      LTMs sent                                   : 0
      LTMs received                               : 0
      LTRs sent                                   : 0
      LTRs received                               : 0
      Sequence number of next LTM request         : 0
      1DMs sent                                   : 0
      Valid 1DMs received                         : 0
      Invalid 1DMs received                       : 0
      DMMs sent                                   : 0
      DMRs sent                                   : 0
      Valid DMRs received                         : 0
      Invalid DMRs received                       : 0
    Remote MEP count: 1
      Identifier     MAC address        State    Interface
        200      00:05:85:73:39:4a       ok    xe-5/0/0.0
```

Configuring Remote MEP Action Profile Support

Based on values of **interface-status-tlv** and **port-status-tlv** in the received CCM packets, a specific action, such as **interface-down**, can be taken using the **action-profile** options.

Multiple action profiles can be configured on the router, but only one action profile can be assigned to a remote MEP.

The action profile can be configured with at least one event to trigger the action; but the action will be triggered if any one of these events occurs. It is not necessary for all of the configured events to occur to trigger **action**.

An action-profile can be applied only at the remote MEP level.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile tlv-action {
  event {
    # If interface status tlv with value specified in the config is received
    interface-status-tlv down|lower-layer-down;
    # If port status tlv with value specified in the config is received
    port-status-tlv blocked;
    # If connectivity is lost to the peer */
    adjacency-loss;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
  default-actions interface-down;
  }
# domains
  maintenance-domain identifier {
    # maintenance domain level (0-7)
    level number;
    # association
    maintenance-association identifier {
      mep identifier {
        interface ge-x/y/z.w;
        remote-mep identifier {
          # Apply the action-profile for the remote MEP
          action-profile tlv-action;
        }
      }
    }
  }
```

Monitoring a Remote MEP Action Profile

You can use the **show oam ethernet connectivity-fault-management mep-database** command to view the action profile status of a remote MEP, as in the following example:

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
```

 Copyright © 2013, Juniper Networks, Inc.

```
 Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
 MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
 Auto-discovery: enabled, Priority: 0
 Interface status TLV: none, Port status TLV: none  # last status TLVs transmitted
by the router
 Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

 Remote MEP identifier: 200, State: ok # displays the remote MEP name and state

   MAC address: 00:05:85:73:96:1f, Type: Configured
   Interface: ge-1/0/8.0
   Last flapped: Never
   Remote defect indication: false
   Port status TLV: none
   Interface status TLV: lower-layer-down
   Action profile: juniper  # displays remote MEP’s action profile identifier
     Last event: Interface-status-tlv lower-layer-down  # last remote MEP event

                                                         # to trigger action
     Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
                                # action occurrence time
```

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

Configuring MAC Flush Message Processing in CET Mode

In carrier Ethernet transport (CET) mode, MX Series routers are used as provider edge (PE) routers, and Nokia Siemens Networks A2200 Carrier Ethernet Switches (referred to as E-domain devices) that run standard-based protocols are used in the access side. On the MX Series routers, VPLS pseudowires are configured dynamically through label distribution protocol (LDP). On the E-domain devices, topology changes are detected through connectivity fault management (CFM) sessions running between the E-domain devices and the MX Series PE routers. The MX Series PE routers can bring the carrier Ethernet interface down if there is CFM connectivity loss. This triggers a local MAC flush as well as a targeted label distribution protocol (T-LDP) MAC flush notification that gets sent towards the remote MX Series PEs to trigger MAC flush on them.

In CET inter-op mode, MX Series routers need to interoperate with the Nokia Siemens Networks Ax100 Carrier Ethernet access devices (referred to as A-domain devices) that run legacy protocols. Nokia Siemens Networks A4100 and A8100 devices act as an intermediate between the MX Series PE routers and A-domain devices. These intermediate devices perform interworking function (IWF) procedures so that operations administration management (OAM) sessions can be run between MX Series routers and A-domain devices. There are no VPLS pseudowires between the MX Series PE routers and the Nokia Siemens Networks A4100 and A8100 intermediate devices, so there is no LDP protocol running between the PE routers to send topology change notifications. In order to communicate topology changes, MX Series routers can trigger a MAC flush and propagate it in the core. MX Series routers can use action profiles based upon the connection protection type length value (TLV) event. The action profile brings down the carrier edge logical interface in MX Series PE routers, which will trigger a local MAC flush and also propagate the topology change to the core using LDP notification.

For VPLS there is no end-to-end connectivity monitored. The access rings are independently monitored by running CFM down multiple end points (MEPs) on the working and protection paths for each of the services between the E-domain devices and the MX Series PE routers, and between the A-domain devices and the MX Series PE routers the IWF hosted by the Nokia Siemens Networks A-4100 devices. When there is a connectivity failure on the working path, the Nokia Siemens Networks Ax200 devices perform a switchover to the protection path, triggering a topology change notification (in the form of TLVs carried in CCM) to be sent on the active path.

 Copyright © 2013, Juniper Networks, Inc.

Figure 26: CET inter-op Dual Homed Topology

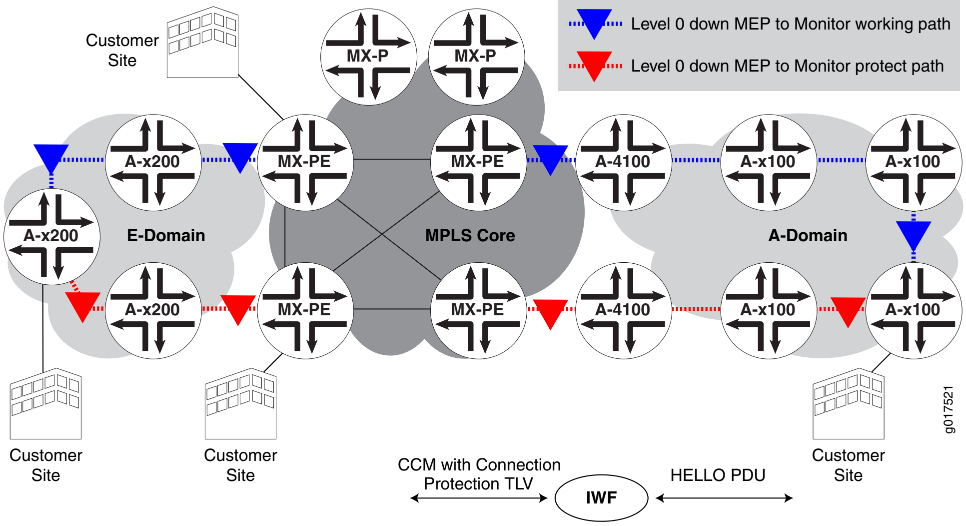

Figure 26 on page 289 describes the dual homed topology on MX Series PE routers connected to the A-domain. When an A-domain device triggers a switchover, it starts switching the service traffic to the new active path. This change is communicated in the HELLO protocol data units (PDUs) sent by that A-domain device on the working and protection paths. When the IWF in A4100 recieves these HELLO PDUs, it converts them to standard CCM messages and also inserts a connection protection TLV. The “Protection-in-use” field of the connection protection TLV is encoded with the currently active path, and is included in the CCM message. CCM messages are received by the MX Series PE routers through the VLAN spoke in A4100. In the above dual homed scenario, one MX Series PE router monitors the working path, and the other MX Series PE router monitors the protection path.

A MAC flush occurs when the CFM session that is monitoring the working path detects that the service traffic has moved to the protection path or when the CFM session that is monitoring the protection path detects that the service traffic has moved to the working path.

Figure 27: CET inter-op Dual Attached Topology

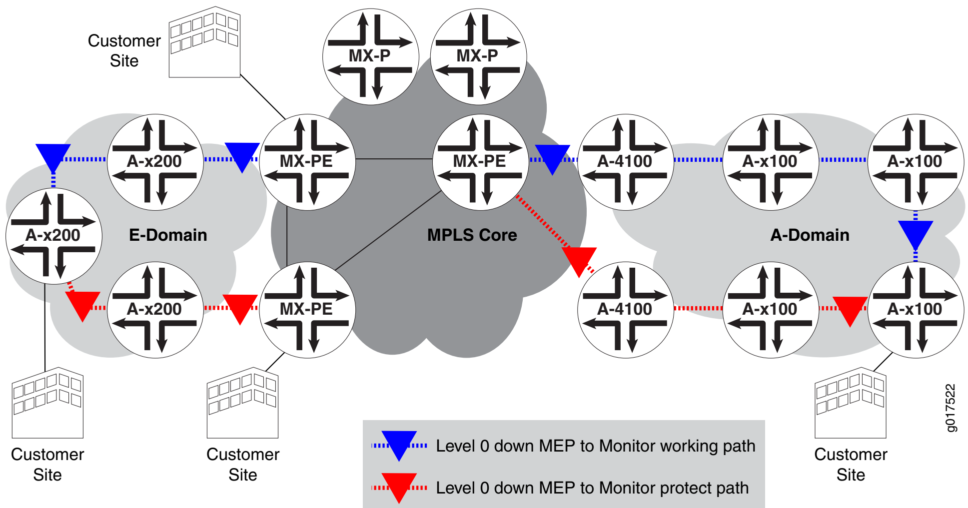

Figure 27 on page 290 describes the dual attached topology on MX Series PE routers connected to the A-domain. The MAC flush mechanism used in this case is also the same as the one used for the A-domain in the dual homed scenario (Figure 1). However in this case both the CFM sessions are hosted by only one MX Series PE router. When Ax100 in the A-domain detects topology changes, the MX Series PE router receives the connection protection TLV in the CCM message for the working and protection paths with the value of "Protection-in-use" indicating which path is the active one. Based upon the event that is generated for the CFM session, the MX Series PE router will bring down the appropriate interface which will trigger a local MAC flush.

Configuring a Connection Protection TLV Action Profile

An action profile can be configured to perform the **interface-down** action based on the values of **connection-protection-tlv** in the received CCM packets.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
    # If a connection protection TLV with a "Protection-in-use" value of SET is received */
    connection-protection-tlv <using-protection-path>;
    # If a connection protection TLV with a "Protection-in-use" value of RESET is received
      */
    connection-protection-tlv <using-working-path>;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
}
```

Related Documentation

- *connection-protection-tlv*

 Copyright © 2013, Juniper Networks, Inc.

Configuring M120 and MX Series Routers for CCC Encapsulated Packets

IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets Overview

Layer 2 virtual private network (L2VPN) is a type of virtual private network service used to transport customer's private Layer 2 traffic (for example, Ethernet, ATM or Frame Relay) over the service provider's shared IP/MPLS infrastructure. The service provider edge (PE) router must have an interface with circuit cross-connect (CCC) encapsulation to switch the customer edge (CE) traffic to the public network.

The IEEE 802.1ag Ethernet Connectivity Fault Management (CFM) is an OAM standard used to perform fault detection, isolation, and verification on virtual bridge LANs. M120 and MX Series routers provide CFM support for bridge/VPLS/routed interfaces and support 802.1ag Ethernet OAM for CCC encapsulated packets.

CFM Features Supported on Layer 2 VPN Circuits

CFM features supported on L2VPN circuits are as follows:

- Creation of up/down MEPs at any level on the CE-facing logical interfaces.
- Creation of MIPs at any level on the CE-facing logical interfaces.
- Support for continuity check, loopback, and linkrace protocol.

Copyright © 2013, Juniper Networks, Inc.

- Support for the Y1731 Ethernet Delay measurement protocol.
- Support for action profiles to bring the CE-facing logical interfaces down when loss of connectivity is detected.

Figure 28: Layer 2 VPN Topology

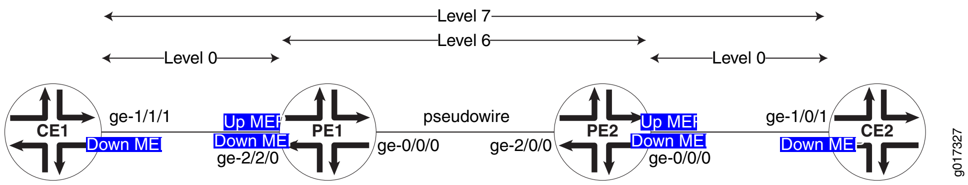

To monitor the L2VPN circuit, a CFM up MEP (Level 6 in Figure 28 on page 292) can be configured on the CE-facing logical interfaces of provider edge routers PE1 and PE2. To monitor the CE-PE attachment circuit, a CFM down MEP can be configured on the customer logical interfaces of CE1-PE1 and CE2-PE2 (Level 0 in Figure 28 on page 292).

Configuring CFM for CCC Encapsulated Packets

The only change from the existing CLI configuration is the introduction of a new command to create a MIP on the CE-facing interface of the PE router.

```
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                # Define a maintenance domains for each default level.
                #; These names are specified as DEFAULT_level_number
                maintenance-domain DEFAULT_x {
                    # L2VPN CE interface
                    interface (ge | xe)-fpc/pic/port.domain;
                }
                {
                level number;
                maintenance-association identifier {
                    mep mep-id {
                        direction (up | down);
                        # L2 VPN CE interface on which encapsulation family CCC is configured.
                        interface (ge | xe)-fpc/pic/port.domain;
                        auto-discovery;
                        priority number;
                    }
                }
            }
        }
    }
    }
}
```

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252

 Copyright © 2013, Juniper Networks, Inc.

Configuring Rate Limiting of Ethernet OAM Messages

M Series, M320 with Enhanced III FPC, M120, M7i and M10 with CFEB, and MX Series routers support rate limiting of Ethernet OAM messages. Depending on the connectivity fault management (CFM) configuration, CFM packets are discarded, sent to the CPU for processing, or flooded to other bridge interfaces. This feature allows the router to intercept incoming CFM packets for prevention of DoS attacks.

You can apply rate limiting of Ethernet OAM messages at either of two CFM policing levels, as follows:

- Global-level CFM policing—uses a policer at the global level to police the CFM traffic belonging to all the sessions.
- Session-level CFM policing—uses a policer created to police the CFM traffic belonging to one session.

To configure global-level CFM policing, include the **policer** statement and its options at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

To configure session-level CFM policing, include the **policer** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *name* level *number* maintenance-association *name*]** hierarchy level.

The following example shows a CFM policer used for rate-limiting CFM:

```
[edit]
firewall {
   policer cfm-policer {
      if-exceeding {
         bandwidth-limit 8k;
         burst-size-limit 2k;
      }
```

Copyright © 2013, Juniper Networks, Inc.

```
        then discard;
    }
}
```

Case 1: Global-Level CFM Policing

This example shows a global level policer, at the CFM level, for rate-limiting CFM. The **continuity-check** ***cfm-policer*** statement at the global **connectivity-fault-management policer** hierarchy level specifies the policer to use for policing all continuity check packets of the CFM traffic belonging to all sessions. The **other** ***cfm-policer1*** statement at the **connectivity-fault-management policer** hierarchy level specifies the policer to use for policing all non-continuity check packets of the CFM traffic belonging to all sessions. The **all** ***cfm-policer2*** statement specifies to police all CFM packets with the specified policer *cfm-policer2*. If the **all** ***policer-name*** option is used, then the user cannot specify the previous **continuity-check** and **other** options.

```
[edit protocols oam ethernet]
connectivity-fault-management {
    policer {
        continuity-check cfm-policer;
        other cfm-policer1 ;
        # all cfm-policer2;
        }
    }
```

Case 2: Session-Level CFM Policing

This example shows a session-level CFM policer used for rate-limiting CFM. The **policer** statement at the session **connectivity-fault-management maintenance-domain** ***md*** **maintenance-association** ***ma*** hierarchy level specifies the policer to use for policing only continuity check packets of the CFM traffic belonging to the specified session. The **other** ***cfm-policer1*** statement at the **connectivity-fault-management maintenance-domain** ***md*** **maintenance-association** ***ma*** hierarchy level specifies the policer to use for policing all non-continuity check packets of the CFM traffic belonging to this session only. The **all** ***cfm-policer2*** statement specifies to police all CFM packets with the specified policer *cfm-policer2*. If the **all** ***policer-name*** option is used, then the user cannot specify the previous **continuity-check** and **other** options.

```
[edit protocols oam ethernet]
connectivity-fault-management {
    maintenance-domain md {
        level number;
        maintenance-association ma {
            continuity-check {
                interval 1s;
            }
            policer {
                continuity-check cfm-policer;
                other cfm-policer1;
                # all cfm-policer2;
                }
                mep 1 {
                    interface ge-3/3/0.0;
                    direction up;
                    auto-discovery;
                }
            }
        }
    }
```

 Copyright © 2013, Juniper Networks, Inc.

In the case of global CFM policing, the same policer is shared across multiple CFM sessions. In per-session CFM policing, a separate policer must be created to rate-limit packets specific to that session.

NOTE:
Service-level policer configuration for any two CFM sessions on the same interface at different levels must satisfy the following constraints if the direction of the sessions is the same:

- If one session is configured with **policer all**, then the other session cannot have a **policer all** or **policer other** configuration.
- If one session is configured with **policer other**, then the other session cannot have a **policer all** or **policer other** configuration.

A commit error will occur if such a configuration is committed.

NOTE: Policers with PBB and MIPs are not supported.

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring 802.1ag Ethernet OAM for VPLS on page 296
- *Ethernet Interfaces*

Configuring 802.1ag Ethernet OAM for VPLS

BEST PRACTICE: The logical interfaces in a VPLS routing instance may have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. VLAN normalization is effectively VLAN translation wherein the VLAN tags of the received packet need to be translated if they are different than the normalized VLAN tags. Configuration is described starting in "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 249 and you should further observe the additional requirements described in this section.

For MX Series routers, the normalized VLAN is specified using one of the following configuration statements in the VPLS routing instance:

- **vlan-id** *vlan-number*
- **vlan-id none**
- **vlan-tags outer** *outer-vlan-number* **inner** *inner-vlan-number*

You must configure **vlan-maps** explicitly on all interfaces belonging to the routing instance.

The following forwarding path considerations must be observed:

- Packet receive path:
  - This is the forwarding path for packets received on the interfaces.
  - 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.
- Packet transmit path:
  - The JUNOS Software uses the router's hardware-based forwarding for CPU-generated packets.
  - For Down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - For Up MEPs, the packet must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packet to be forwarded with this flood route.
  - The router also uses implicit-based forwarding for CPU generated packets. The result is for the flood next hop tied to the flood route to be tied to the filter term. The filter term uses match criteria to correctly identify the host- generated packets.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *connectivity-fault-management*
- *Example: Configuring Ethernet CFM over VPLS*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- Creating the Maintenance Domain on page 252
- Configuring Maintenance Intermediate Points on page 253
- Creating a Maintenance Association on page 255
- Continuity Check Protocol on page 256
- Configuring a Maintenance Endpoint on page 259
- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Linktrace Protocol in CFM on page 267
- Configuring Ethernet Local Management Interface on page 268
- Configuring Port Status TLV and Interface Status TLV on page 275
- Configuring MAC Flush Message Processing in CET Mode on page 288
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 291
- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- *Ethernet Interfaces*

Configuring Unified ISSU for 802.1ag CFM

A unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is automatically enabled for the Connectivity Fault Management (CFM) protocols and interoperates between local and remote maintenance endpoints (MEPs).

The Junos OS provides support for unified ISSU using the loss threshold type length value (TLV), which is automatically enabled for CFM. TLVs are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and optional information in a protocol data unit (PDU). The loss threshold TLV indicates the loss threshold value of a remote MEP. The loss threshold TLV is transmitted as part of the CFM continuity check messages.

NOTE: **Configuring ISSU with CFM (802.1ag) is supported only between two MX routers that support TLV. Interoperation with other vendors is not supported.**

During a unified ISSU, the control plane may go down for several seconds and cause CFM continuity check packets to get dropped. This may cause the remote MEP to detect a

Copyright © 2013, Juniper Networks, Inc.

connectivity loss and mark the MEP as down. To keep the MEP active during a unified ISSU, the loss threshold TLV communicates the minimum threshold value the receiving MEP requires to keep the MEP active. The receiving MEP parses the TLV and updates the loss threshold value, but only if the new threshold value is greater than the locally configured threshold value.

An overview of CFM is described starting in "IEEE 802.1ag OAM Connectivity Fault Management Overview" on page 249, and you should further observe the additional requirements described in this topic.

Table 24 on page 298 shows the Loss Threshold TLV format.

Table 24: Loss Threshold TLV Format

Parameter	Octet (sequence)	Description
Type=31	1	Required. Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.
Length=12	2	Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field.
OUI	3	Optional. Organization unique identifier (OUI), which is controlled by the IEEE and is typically the first three bytes of a MAC address (Juniper OUI 0x009069).
Subtype	1	Optional. Organizationally defined subtype.
Value	4	Optional. Loss threshold value.
Flag	4	Optional. Bit0 (identifies an ISSU is in progress) Bit1-31 (reserved)

Junos OS provides configuration support for the **convey-loss-threshold** statement, allowing you to control the transmission of the loss threshold TLV in continuity check messages PDUs. The **convey-loss-threshold** statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the **convey-loss-threshold** statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress. The Junos OS provides this configuration at the continuity-check level. By default, continuity check messages do not include the loss threshold TLV.

To configure the convey loss threshold, use the **convey-loss-threshold** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check]** hierarchy level.

For the remote MEP, the loss threshold TLV is transmitted only during the unified ISSU if the **convey-loss-threshold** statement is not configured. The remote MEP switches back to the default loss threshold if no loss threshold TLV is received or the TLV has a default threshold value of 3.

Copyright © 2013, Juniper Networks, Inc.

An example of the ISSU configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
        level number;
          maintenance-association identifier {
            continuity-check {
              convey-loss-threshold;
              interval number;
              loss-threshold number;
              hold-interval number;
            }
          }
        }
      }
    }
  }
}
```

The Junos OS saves the last received loss threshold TLV from the remote MEP. You can display the last saved loss threshold TLV that is received by the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***identifier*** **maintenance-association** ***identifier*** **local-mep** ***identifier*** **remote-mep** ***identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md3 maintenance-association ma5 local-mep 2 remote-mep 1
Maintenance domain name: md3, Format: string, Level: 3
  Maintenance association name: ma3, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
    Loss Threshold TLV:
      Loss Threshold: 3  , Flag:    0x0

 Remote MEP identifier: 1, State: ok
    MAC address: 00:1f:12:b7:ce:79, Type: Learned
    Interface: xe-4/1/1.0
    Last flapped: Never
    Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
    Effective loss threshold: 3 frames
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
    Connection Protection TLV:
      Prefer me: no, Protection in use: no, FRR Flag: no
    Loss Threshold TLV:    #Displays last received value
      Loss Threshold: 3  , Flag:    0x0
```

The Junos OS saves the last transmitted loss threshold TLV from a local MEP. You can display the last transmitted loss threshold TLV and the effective loss (operational) threshold for the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** *identifier* **maintenance-association** *identifier* **local-mep** *identifier* **remote-mep** *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md3 maintenance-association ma5 local-mep 2 remote-mep 1
Maintenance domain name: md3, Format: string, Level: 3
  Maintenance association name: ma3, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Connection Protection TLV: yes
    Prefer me: no, Protection in use: no, FRR Flag: no
  Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
    Loss Threshold TLV:    #Displays last transmitted value
      Loss Threshold: 3  , Flag:    0x0

   Remote MEP identifier: 1, State: ok
    MAC address: 00:1f:12:b7:ce:79, Type: Learned
    Interface: xe-4/1/1.0
    Last flapped: Never
    Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
    Effective loss threshold: 3 frames    #Displays operational threshold
    Remote defect indication: falsePort status TLV: none
    Interface status TLV: none
    Connection Protection TLV:
      Prefer me: no, Protection in use: no, FRR Flag: no
    Loss Threshold TLV:
      Loss Threshold: 3  , Flag:    0x0
```

Related Documentation

- *Example: Configuring Ethernet CFM over VPLS*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249

Configuring CCM for Better Scalability

This topic describes how to configure CCM for better scalability. Junos OS provides enhancements to trigger faster protection-switching and convergence in the event of failures in Ethernet domains for Carrier Ethernet services. These enhancements can be used when CE devices in the Ethernet domain detect faster service failures and propagates the information in the interface-status TLV of the continuity-check messages (CCMs). When CCMs are received, PE devices can perform certain actions which facilitates faster protection-switching and convergence.

To configure CCM for better scalability:

- You can apply an action profile to provide faster protection switching for point-to-point network topologies with local switching configured. See *Configuring Faster Protection Switching*.
- You can apply an action profile to provide faster convergence for dual-homed multipoint-to-multipoint network topologies. See *Configuring Faster Convergence*.

 Copyright © 2013, Juniper Networks, Inc.

- You can assign a primary virtual LAN (VLAN) ID in the maintenance association for increased flexibility in the number of tags. See *Configuring a Primary VLAN ID for Increased Flexibility*.
- You can configure a maintenance association to accept a different maintenance association identifier (ID) from a neighbor by including a **remote-maintenance-association** statement. See *Configuring a Remote Maintenance Association to Accept Different ID*.

Related Documentation

- *Configuring Faster Protection Switching*
- *Configuring Faster Convergence*
- *Configuring a Primary VLAN ID for Increased Flexibility*
- *Configuring a Remote Maintenance Association to Accept Different ID*

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 14

Configuring ITU-T Y.1731 Ethernet Service OAM

Service-Level Agreement Measurement

Service-level agreement (SLA) measurement is the process of monitoring the bandwidth, delay, delay variation (jitter), continuity, and availability of a service (E-Line or E-LAN). It enables you to identify network problems before customers are impacted by network defects.

NOTE:
The Ethernet VPN services can be classified into:

- Peer-to-peer-services (E-Line services)—The E-Line services are offered using MPLS-based Layer 2 VPN virtual private wire service (VPWS).
- Multipoint-to-multipoint services (E-LAN services)—The E-LAN services are offered using MPLS-based virtual private LAN service (VPLS).

For more information, see the *Junos VPNs Configuration Guide*.

In Junos OS, SLA measurements are classified into:

- On-demand mode—In on-demand mode, the measurements are triggered through the CLI. For more information, see "On-Demand Mode" on page 311.
- Proactive mode—In proactive mode, the measurements are triggered by an iterator application. For more information, see "Proactive Mode" on page 312.

For more information about frame delay measurement, see "Ethernet Frame Delay Measurements Overview" on page 304. For more information about frame loss measurement, see "Ethernet Frame Loss Measurement Overview" on page 310. Note that Ethernet frame delay measurement and Ethernet frame loss measurement are not supported on the **ae** interface.

Related Documentation

- Proactive Mode on page 312.
- On-Demand Mode on page 311.
- *Ethernet Interfaces*

Ethernet Frame Delay Measurements Overview

 Copyright © 2013, Juniper Networks, Inc.

ITU-T Y.1731 Frame Delay Measurement Feature

The IEEE 802.3-2005 standard for Ethernet Operations, Administration, and Maintenance (OAM) defines a set of link fault management mechanisms to detect and report link faults on a single point-to-point Ethernet LAN.

Junos OS supports key OAM standards that provide for automated end-to-end management and monitoring of Ethernet service by service providers:

- *IEEE Standard 802.1ag*, also known as "Connectivity Fault Management (CFM)."
- *ITU-T Recommendation Y.1731*, which uses different terminology than IEEE 802.1ag and defines Ethernet service OAM features for fault monitoring, diagnostics, and performance monitoring.

These capabilities allow operators to offer binding service-level agreements (SLAs) and generate new revenues from rate- and performance-guaranteed service packages that are tailored to the specific needs of their customers.

Ethernet CFM

The IEEE 802.1ag standard for connectivity fault management (CFM) defines mechanisms to provide for end-to-end Ethernet service assurance over any path, whether a single link or multiple links spanning networks composed of multiple LANs.

For Ethernet interfaces on M320, MX Series, and T Series routers, Junos OS supports the following key elements of the Ethernet CFM standard:

- Fault monitoring using the IEEE 802.1ag Ethernet OAM Continuity Check protocol
- Path discovery and fault verification using the IEEE 802.1ag Ethernet OAM Linktrace protocol
- Fault isolation using the IEEE 802.1ag Ethernet OAM Loopback protocol

In a CFM environment, network entities such as network operators, service providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management and end-to-end monitoring, and still avoid security breaches.

Figure 29 on page 306 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Copyright © 2013, Juniper Networks, Inc.

Figure 29: Relationship of MEPs, MIPs, and Maintenance Domain Levels

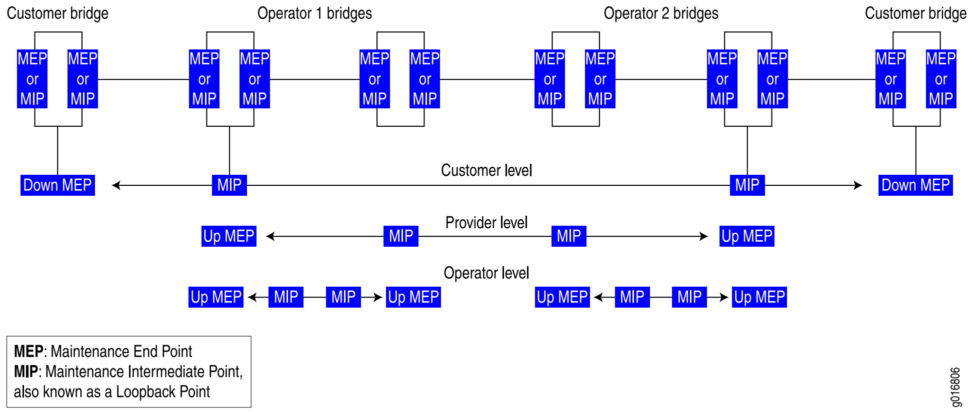

NOTE: Maintenance intermediate points (MIP) are not supported on the ACX Series routers.

Ethernet Frame Delay Measurement

Two key objectives of OAM functionality are to measure quality-of-service attributes such as frame delay and frame delay variation (also known as “frame jitter”). Such measurements can enable you to identify network problems before customers are impacted by network defects.

Junos OS supports Ethernet frame delay measurement between MEPs configured on Ethernet physical or logical interfaces on MX Series routers. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor SLAs. Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. The Junos OS implementation of Ethernet frame delay measurement (ETH-DM) is fully compliant with the ITU-T Recommendation Y.1731, *OAM Functions and Mechanisms for Ethernet-based Networks*. The recommendation defines OAM mechanisms for operating and maintaining the network at the Ethernet service layer, which is called the "ETH layer" in ITU-T terminology.

MX Series routers with modular port concentrators (MPCs) and 10-Gigabit Ethernet MPCs with SFP+ support ITU-T Y.1731 functionality on VPLS for frame-delay and delay-variation.

One-Way Ethernet Frame Delay Measurement

In one-way ETH-DM mode, a series of frame delay and frame delay variation values are calculated based on the time elapsed between the time a measurement frame is sent from the initiator MEP at one router and the time when the frame is received at the receiver MEP at the other router.

Copyright © 2013, Juniper Networks, Inc.

1DM Transmission

When you start a one-way frame delay measurement, the router sends 1DM frames—frames that carry the protocol data unit (PDU) for a one-way delay measurement—from the initiator MEP to the receiver MEP at the rate and for the number of frames you specify. The router marks each 1DM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

1DM Reception

When an MEP receives a 1DM frame, the router that contains the receiver MEP measures the one-way delay for that frame (the difference between the time the frame was received and the timestamp contained in the frame itself) and the delay variation (the difference between the current and previous delay values).

One-Way ETH-DM Statistics

The router that contains the receiver MEP stores each set of one-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

One-Way ETH-DM Frame Counts

Each router counts the number of one-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number of 1DM frames sent.
- For a receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

Synchronization of System Clocks

The accuracy of one-way delay calculations depends on close synchronization of the system clocks at the initiator MEP and receiver MEP.

The accuracy of one-way delay variation is not dependent on system clock synchronization. Because delay variation is simply the difference between consecutive one-way delay values, the out-of-phase period is eliminated from the frame jitter values.

NOTE: For a given one-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only on the router that contains the receiver MEP.

Two-Way Ethernet Frame Delay Measurement

In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

DMM Transmission

When you start a two-way frame delay measurement, the router sends delay measurement message (DMM) frames— frames that carry the PDU for a two-way ETH-DM request—from the initiator MEP to the responder MEP at the rate and for the number of frames you specify. The router marks each DMM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

DMR Transmission

When an MEP receives a DMM frame, the responder MEP responds with a delay measurement reply (DMR) frame, which carries ETH-DM reply information and a copy of the timestamp contained in the DMM frame.

DMR Reception

When an MEP receives a valid DMR, the router that contains the MEP measures the two-way delay for that frame based on the following sequence of timestamps:

1. TI_{TxDMM}
2. TR_{RxDMM}
3. TR_{TxDMR}
4. TI_{RxDMR}

A two-way frame delay is calculated as follows:

$$[TI_{RxDMR} - TI_{TxDMM}] - [TR_{TxDMR} - TR_{RxDMM}]$$

The calculation show that frame delay is the difference between the time at which the initiator MEP sends a DMM frame and the time at which the initiator MEP receives the associated DMR frame from the responder MEP, minus the time elapsed at the responder MEP.

The delay variation is the difference between the current and previous delay values.

Two-Way ETH-DM Statistics

The router that contains the initiator MEP stores each set of two-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

 Copyright © 2013, Juniper Networks, Inc.

Two-Way ETH-DM Frame Counts

Each router counts the number of two-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number DMM frames transmitted, the number of valid DMR frames received, and the number of invalid DMR frames received.
- For a responder MEP, the router counts the number of DMR frames sent.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

NOTE: **For a given two-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only at the router that contains the initiator MEP.**

Choosing Between One-Way and Two-Way ETH-DM

One-way frame delay measurement requires that the system clocks at the initiator MEP and receiver MEP are closely synchronized. Two-way frame delay measurement does not require synchronization of the two systems. If it is not practical for the clocks to be synchronized, two-way frame delay measurements are more accurate.

When two systems are physically close to each other, their one-way delay values are very high compared to their two-way delay values. One-way delay measurement requires that the timing for the two systems be synchronized at a very granular level, and MX Series routers currently do not support this granular synchronization.

Restrictions for Ethernet Frame Delay Measurement

The following restrictions apply to the Ethernet frame delay measurement feature:

- The ETH-DM feature is not supported on aggregated Ethernet interfaces or label-switched interface. (LSI) pseudowires.
- Hardware-assisted timestamping for ETH-DM frames in the reception path is only supported for MEP interfaces on Enhanced DPCs and Enhanced Queuing DPCs in MX Series routers. For information about hardware-assisted timestamping, see “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 326 and “Enabling the Hardware-Assisted Timestamping Option” on page 336.
- Ethernet frame delay measurements can be triggered only when the distributed periodic packet management daemon (**ppm**) is enabled. For more information about this limitation, see “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 326 and “Ensuring That Distributed ppm Is Not Disabled” on page 335.
- You can monitor only one session at a time to the same remote MEP or MAC address. For more information about starting an ETH-DM session, see “Starting an ETH-DM Session” on page 338.

Copyright © 2013, Juniper Networks, Inc.

- ETH-DM statistics are collected at only one of the two peer routers in the ETH-DM session. For a one-way ETH-DM session, you can display frame ETH-DM statistics at the receiver MEP only, using ETH-DM-specific **show** commands. For a two-way ETH-DM session, you can display frame delay statistics at the initiator MEP only, using the same ETH-DM-specific **show** commands. For more information, see "Managing ETH-DM Statistics and ETH-DM Frame Counts" on page 349.
- ETH-DM frame counts are collected at both MEPs and are stored in the respective CFM databases.
- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. Therefore, the collection of ETH-DM statistics and ETH-DM frame counters has to be restarted, after the switchover is complete. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *Junos OS High Availability Library for Routing Devices*.
- Accuracy of frame delay statistics is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Related Documentation

- Ethernet Frame Loss Measurement Overview on page 310
- Example: One-Way Ethernet Frame Delay Measurement on page 358
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 326
- Guidelines for Starting an ETH-DM Session on page 328
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 330
- On-Demand Mode on page 311
- Proactive Mode on page 312
- *Ethernet Interfaces*

Ethernet Frame Loss Measurement Overview

The key objectives of the OAM functionality are to measure quality-of-service attributes such as frame delay, frame delay variation (also known as "frame jitter"), and frame loss. Such measurements enable you to identify network problems before customers are impacted by network defects. For more information about Ethernet frame delay measurement, see "Ethernet Frame Delay Measurements Overview" on page 304.

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on MX Series routers and is presently supported only for VPWS service. ETH-LM is used by operators to collect counter values applicable for ingress and egress service frames. These counters maintain a count of transmitted and received data frames between a pair of MEPs. Ethernet frame loss measurement is performed by sending frames with ETH-LM information to a peer MEP and similarly receiving frames with ETH-LM information

 Copyright © 2013, Juniper Networks, Inc.

from the peer MEP. This type of frame loss measurement is also known as single-ended Ethernet loss measurement.

ETH-LM supports the following frame loss measurements:

- Near-end frame loss measurement—Measurement of frame loss associated with ingress data frames.
- Far-end frame loss measurement—Measurement of frame loss associated with egress data frames.

NOTE: **The proactive and dual-ended loss measurement functionality of ITU-T Y1731 is not supported on the ACX Series routers.**

The Junos OS implementation of Ethernet frame delay measurement (ETH-DM) is fully compliant with the ITU-T Recommendation Y.1731, as described in *OAM Functions and Mechanisms for Ethernet-Based Networks*. The recommendation defines OAM mechanisms for operating and maintaining the network at the Ethernet service layer, which is called the "ETH layer" in ITU-T terminology.

Related Documentation

- Managing Continuity Measurement Statistics on page 358
- On-Demand Mode on page 311
- Proactive Mode on page 312
- *Ethernet Interfaces*

On-Demand Mode

In on-demand mode, the measurements are triggered by the user through the CLI.

When the user triggers the delay measurement through the CLI, the delay measurement request that is generated is as per the frame formats specified by the ITU-T Y.1731 standard. For two-way delay measurement, the server-side processing can be delegated to the Packet Forwarding Engine to prevent overloading on the Routing Engine. For more information, see “Configuring Routers to Support an ETH-DM Session” on page 334. When the server-side processing is delegated to the Packet Forwarding Engine, the delay measurement message (DMM) frame **receive** counters and delay measurement reply (DMR) frame **transmit** counters are not displayed by the **show** command.

When the user triggers the loss measurement through the CLI, the router sends the packets in standard format along with the loss measurement TLV. By default, the **session-id-tlv** argument is included in the packet to allow concurrent loss measurement sessions from same local MEP. You can also disable the session ID TLV by using the **no-session-id-tlv** argument.

Single-ended ETH-LM is used for on-demand operation, administration, and maintenance purposes. An MEP sends frames with ETH-LM request information to its peer MEP and receives frames with ETH-LM reply information from its peer MEP to carry out loss

Copyright © 2013, Juniper Networks, Inc.

measurements. The protocol data unit (PDU) used for a single-ended ETH-LM request is referred to as a loss measurement message (LMM) and the PDU used for a single-ended ETH-LM reply is referred to as a loss measurement reply (LMR).

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 304
- Ethernet Frame Loss Measurement Overview on page 310
- Proactive Mode on page 312
- Configuring Routers to Support an ETH-DM Session on page 334.
- *Ethernet Interfaces*

Proactive Mode

In proactive mode, SLA measurements are triggered by an iterator application. An iterator is designed to periodically transmit SLA measurement packets in form of ITU-Y.1731-compliant frames for two-way delay measurement or loss measurement on MX Series routers. This mode differs from on-demand SLA measurement, which is user initiated. The iterator sends periodic delay or loss measurement request packets for each of the connections registered to it. Iterators make sure that measurement cycles do not occur at the same time for the same connection to avoid CPU overload. Junos OS supports proactive mode for VPWS. For an iterator to form a remote adjacency and to become functionally operational, the continuity check message (CCM) must be active between the local and remote MEP configurations of the connectivity fault management (CFM). Any change in the iterator adjacency parameters resets the existing iterator statistics and restarts the iterator. Here, the term adjacency refers to a pairing of two endpoints (either connected directly or virtually) with relevant information for mutual understanding, which is used for subsequent processing. For example, the iterator adjacency refers to the iterator association between the two endpoints of the MEPs.

For every DPC or MPC, only 30 iterator instances for a cycle time value of 10 milliseconds (ms) are supported. In Junos OS, 255 iterator profile configurations and 2000 remote MEP associations are supported.

Iterators with cycle time value less than 100 ms are supported only for infinite iterators, whereas the iterators with cycle time value greater than 100 ms are supported for both finite and infinite iterators. Infinite iterators are iterators that run infinitely until the iterator is disabled or deactivated manually.

A VPWS service configured on a router is monitored for SLA measurements by registering the connection (here, the connection is a pair of remote and local MEPs) on an iterator and then initiating periodic SLA measurement frame transmission on those connections. The end-to-end service is identified through a maintenance association end point (MEP) configured at both ends.

For two-way delay measurement and loss measurement, an iterator sends a request message for the connection in the list (if any) and then sends a request message for the connection that was polled in the former iteration cycle. The back-to-back request messages for the SLA measurement frames and their responses help in computing delay variation and loss measurement.

 Copyright © 2013, Juniper Networks, Inc.

The Y.1731 frame transmission for a service attached to an iterator continues endlessly unless intervened and stopped by an operator or until the iteration-count condition is met. To stop the iterator from sending out any more proactive SLA measurement frames, the operator must perform one of the following tasks:

- Enable the **deactivate sla-iterator-profile** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance association *ma-name* mep *mep-id* remote-mep *mep-id*]** hierarchy level. For more information, see *Verifying the Configuration of an Iterator Profile*.
- Provision a **disable** statement under the corresponding iterator profile at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]** hierarchy level. For more information, see "Configuring an Iterator Profile" on page 323.

Ethernet Delay Measurements and Loss Measurement by Proactive Mode

In two-way delay measurement, the delay measurement message (DMM) frame is triggered through an iterator application. The DMM frame carries an iterator type, length, and value (TLV) in addition to the fields described in standard frame format and the server copies the iterator TLV from the DMM frame to the delay measurement reply (DMR) frame.

In one-way delay variation computation using the two-way delay measurement method, the delay variation computation is based on the timestamps that are present in the DMR frame (and not the 1DM frame). Therefore, there is no need for client-side and server-side clocks to be in sync. Assuming that the difference in their clocks remains constant, the one-way delay variation results are expected to be fairly accurate. This method also eliminates the need to send separate 1DM frames just for the one-way delay variation measurement purpose.

In proactive mode for loss measurement, the router sends packets in standard format along with loss measurement TLV and iterator TLV.

Related Documentation

- Configuring an Iterator Profile on page 323
- Configuring a Remote MEP with an Iterator Profile on page 324
- Ethernet Frame Delay Measurements Overview on page 304
- Ethernet Frame Loss Measurement Overview on page 310
- *Verifying the Configuration of an Iterator Profile*
- Managing Iterator Statistics on page 353
- On-Demand Mode on page 311
- *Ethernet Interfaces*

Ethernet Failure Notification Protocol Overview

The Failure Notification Protocol (FNP) is a failure notification mechanism that detects failures in Point-to-Point Ethernet transport networks on MX Series routers. If a node link fails, FNP detects the failure and sends out FNP messages to the adjacent nodes that a circuit is down. Upon receiving the FNP message, nodes can redirect traffic to the protection circuit.

NOTE: **FNP is supported on E-Line services only.**

An E-Line service provides a secure Point-to-Point Ethernet connectivity between two user network interfaces (UNIs). E-Line services are a protected service and each service has a working circuit and protection circuit. CFM is used to monitor the working and protect paths. CCM intervals result in failover time in hundreds of milliseconds or a few seconds. FNP provides service circuit failure detection and propagation in less than 50ms and provide 50ms failover for E-Line services.

The MX router acts as a PE node and handles the FNP messages received on the management VLAN and the FNP messages received on both the Ethernet interfaces and PWs created for the management VPLS. MX-series routers do not initiate FNP messages and responds only to FNP messages generated by devices in the Ethernet Access network. FNP can be enabled only on logical interfaces that are part of a VPLS routing instance, and no physical interfaces in that VPLS routing instance should have CCM configured. FNP can be enabled only on one logical interface per physical interface.

All E-Line services are configured as layer 2 circuits with edge protection. A VLAN associated with the working circuit or protection circuit must map to a logical interface. No trunk port or access port is supported in the ring link for VLANs used by E-LINE services. FNP does not control the logical interface associated with protection circuit. Only E-Line service whose termination point is not in an MX node is controlled by FNP.

FNP supports graceful restart and the Graceful Routing Engine switchover (GRES) features.

Related Documentation

- Configuring the Failure Notification Protocol on page 365
- *show oam ethernet fnp interface*
- *show oam ethernet fnp status*
- *show oam ethernet fnp messages*
- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Ethernet Synthetic Loss Measurement Overview

Ethernet synthetic loss measurement (ETH-SLM) is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic. Each maintenance association end point (MEP) performs frame loss measurements, which contribute to unavailable time.

A near-end frame loss specifies frame loss associated with ingress data frames and a far-end frame loss specifies frame loss associated with egress data frames. Both near-end and far-end frame loss measurements contribute to near-end severely errored seconds and far-end severely errored seconds that are used in combination to determine unavailable time. ETH-SLM is performed using synthetic loss message (SLM) and synthetic loss reply (SLR) frames. ETH-SLM facilitates each MEP to perform near-end and far-end synthetic frame loss measurements by using synthetic frames because a bidirectional service is defined as unavailable if either of the two directions is determined to be unavailable.

There are the two types of frame loss measurement, defined by the ITU-T Y.1731 standards, ETH-LM and ETH-SLM. Junos OS supports only single-ended ETH-SLM. In single-ended ETH-SLM, each MEP sends frames with the ETH-SLM request information to its peer MEP and receives frames with ETH-SLM reply information from its peer MEP to perform synthetic loss measurements. Single-ended ETH-SLM is used for proactive or on-demand OAM to perform synthetic loss measurements applicable to point-to-point Ethernet connection. This method allows a MEP to initiate and report far-end and near-end loss measurements associated with a pair of MEPs that are part of the same maintenance entity group (MEG).

Single-ended ETH-SLM is used to perform on-demand or proactive tests by initiating a finite amount of ETH-SLM frames to one or multiple MEP peers and receiving the ETH-SLM reply from the peers. The ETH-SLM frames contain the ETH-SLM information that is used to measure and report both near-end and far-end synthetic loss measurements. Service-level agreement (SLA) measurement is the process of monitoring the bandwidth, delay, delay variation (jitter), continuity, and availability of a service. It enables you to identify network problems before customers are impacted by network defects. In proactive mode, SLA measurements are triggered by an iterator application. An iterator is designed to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for synthetic frame loss measurement . This mode differs from on-demand SLA measurement, which is user initiated. In on-demand mode, the measurements are triggered by the user through the CLI. When the user triggers the ETH-SLM through the CLI, the SLM request that is generated is as per the frame formats specified by the ITU-T Y.1731 standard.

Related Documentation

- Transmission of ETH-SLM Messages on page 319
- Format of ETH-SLM Messages on page 317
- Guidelines for Configuring ETH-SLM on page 321
- Scenarios for Configuration of ETH-SLM on page 316

- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Starting an On-Demand ETH-SLM Session on page 343
- Troubleshooting Failures with ETH-SLM on page 348

- *Ethernet Interfaces*

Scenarios for Configuration of ETH-SLM

ETH-SLM measures near-end and far-end frame loss between two MEPs that are part of the same MEG level. You can configure ETH-SLM to measure synthetic loss for both upward-facing or upstream MEP and downward-facing or downstream MEP. This section describes the following scenarios for the operation of ETH-SLM:

Upstream MEP in MPLS Tunnels

Consider a scenario in which a MEP is configured between the user network interfaces (UNIs) of two MX Series routers, MX1 and MX2, in the upstream direction. MX1 and MX2 are connected over an MPLS core network. ETH-SLM measurements are performed between the upstream MEP in the path linking the two routers. Both MX1 and MX2 can initiate on-demand or proactive ETH-SLM, which can measure both far-end and near-end loss at MX1 and MX2, respectively. The two UNIs are connected using MPLS-based Layer 2 VPN virtual private wire service (VPWS).

Downstream MEP in Ethernet Networks

Consider a scenario in which a MEP is configured between two MX Series routers, MX1 and MX2, on the Ethernet interfaces in the downstream direction. MX1 and MX2 are connected in an Ethernet topology and downstream MEP is configured toward the Ethernet network. ETH-SLM measurements are performed between the downstream MEP in the path linking the two routers. ETH-SLM can be measured in the path between these two routers.

Consider another scenario in which a MEP is configured in the downstream direction and service protection for a VPWS over MPLS is enabled by specifying a working path or protect path on the MEP. Service protection provides end-to-end connection protection of the working path in the event of a failure. To configure service protection, you must create two separate transport paths—a working path and a protect path. You can specify the working path and protect path by creating two maintenance associations. To associate the maintenance association with a path, you must configure the MEP interface in the maintenance association and specify the path as working or protect.

In a sample topology, an MX Series router, MX1, is connected to two other MX Series routers, MX2 and MX3, over an MPLS core. The connectivity fault management (CFM) session between MX1 and MX2 is the working path on the MEP and the CFM session between MX1 and MX3 is the protect path on the MEP. MX2 and MX3 are, in turn, connected on Ethernet interfaces to MX4 in the access network. Downstream MEP is configured between MX1 and MX4 that passes through MX2 (working CFM session) and also between MX1 and MX4 that passes through MX3 (protected CFM session). ETH-SLM

 Copyright © 2013, Juniper Networks, Inc.

is performed between these downstream MEPs. In both the downstream MEPs, the configuration is performed on MX1 and MX4 UNIs, similar to upstream MEP.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Transmission of ETH-SLM Messages on page 319
- Format of ETH-SLM Messages on page 317
- Guidelines for Configuring ETH-SLM on page 321
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Starting an On-Demand ETH-SLM Session on page 343
- Troubleshooting Failures with ETH-SLM on page 348
- *Ethernet Interfaces*

Format of ETH-SLM Messages

Synthetic loss messages (SLMs) support single-ended Ethernet synthetic loss measurement (ETH-SLM) requests. This topic contains the following sections that describe the formats of the SLM protocol data units (PDUs), SLR PDUs, and the data iterator type length value (TLV).

SLM PDU Format

The SLM PDU format is used by a MEP to transmit SLM information. The following components are contained in SLM PDUs:

- Source MEP ID—Source MEP ID is a 2-octet field where the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.
- Test ID—Test ID is a 4-octet field set by the transmitting MEP and is used to identify a test when multiple tests run simultaneously between MEPs (including both concurrent on-demand and proactive tests).
- TxFCf—TxFCf is a 4-octet field that carries the number of SLM frames transmitted by the MEP toward its peer MEP.

The following are the fields in an SLM PDU:

- MEG Level—Configured maintenance domain level in the range 0–7.
- Version—0.
- OpCode—Identifies an OAM PDU type. For SLM, it is 55.
- Flags—Set to all zeros.
- TLV Offset—16.

Copyright © 2013, Juniper Networks, Inc.

- Source MEP ID—A 2-octet field used to identify the MEP transmitting the SLM frame. In this 2-octet field, the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.
- RESV—Reserved fields are set to all zeros.
- Test ID—A 4-octet field set by the transmitting MEP and used to identify a test when multiple tests run simultaneously between MEPs (including both concurrent on-demand and proactive tests).
- TxFCf—A 4-octet field that carries the number of SLM frames transmitted by the MEP toward its peer MEP.
- Optional TLV—A data TLV may be included in any SLM transmitted. For the purpose of ETH-SLM, the value part of data TLV is unspecified.
- End TLV—All zeros octet value.

SLR PDU Format

The synthetic loss reply (SLR) PDU format is used by a MEP to transmit SLR information. The following are the fields in an SLR PDU:

- MEG Level—A 3-bit field the value of which is copied from the last received SLM PDU.
- Version—A 5-bit field the value of which is copied from the last received SLM PDU.
- OpCode—Identifies an OAM PDU type. For SLR, it is set as 54.
- Flags—A 1-octet field copied from the SLM PDU.
- TLV Offset—A 1-octet field copied from the SLM PDU.
- Source MEP ID—A 2-octet field copied from the SLM PDU.
- Responder MEP ID—A 2-octet field used to identify the MEP transmitting the SLR frame.
- Test ID—A 4-octet field copied from the SLM PDU.
- TxFCf—A 4-octet field copied from the SLM PDU.
- TxFCb—A 4 octet field. This value represents the number of SLR frames transmitted for this test ID.
- Optional TLV—The value is copied from the SLM PDU, if present.
- End TLV—A 1-octet field copied from the SLM PDU.

Data Iterator TLV Format

The data iterator TLV specifies the data TLV portion of the Y.1731 data frame. The MEP uses a data TLV when the MEP is configured to measure delay and delay variation for different frame sizes. The following are the fields in a data TLV:

- Type—Identifies the TLV type; value for this TLV type is Data (3).
- Length—Identifies the size, in octets, of the Value field containing the data pattern. The maximum value of the Length field is 1440.

 Copyright © 2013, Juniper Networks, Inc.

- Data pattern—An *n*-octet (*n* denotes length) arbitrary bit pattern. The receiver ignores it.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Transmission of ETH-SLM Messages on page 319
- Guidelines for Configuring ETH-SLM on page 321
- Scenarios for Configuration of ETH-SLM on page 316
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Starting an On-Demand ETH-SLM Session on page 343
- Troubleshooting Failures with ETH-SLM on page 348
- *Ethernet Interfaces*

Transmission of ETH-SLM Messages

The ETH-SLM functionality can process multiple synthetic loss message (SLM) requests simultaneously between a pair of MEPs. The session can be a proactive or an on-demand SLM session. Each SLM request is identified uniquely by a test ID.

A MEP can send SLM requests or respond to SLM requests. A response to an SLM request is called a synthetic loss reply (SLR). After a MEP determines an SLM request by using the test ID, the MEP calculates the far-end and near-end frame loss on the basis of the information in the SLM message or the SLM protocol data unit (PDU).

A MEP maintains the following local counters for each test ID and for each peer MEP being monitored in a maintenance entity for which loss measurements are to be performed:

- TxFCl—Number of synthetic frames transmitted toward the peer MEP for a test ID. A source MEP increments this number for successive transmission of synthetic frames with ETH-SLM request information while a destination or receiving MEP increments this value for successive transmission of synthetic frames with the SLR information.
- RxFCl—Number of synthetic frames received from the peer MEP for a test ID. A source MEP increments this number for successive reception of synthetic frames with SLR information while a destination or receiving MEP increments it for successive reception of synthetic frames with ETH-SLM request information.

The following sections describe the phases of processing of SLM PDUs to determine synthetic frame loss:

Initiation and Transmission of SLM Requests

A MEP periodically transmits an SLM request with the OpCode field set as 55. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the

Copyright © 2013, Juniper Networks, Inc.

local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TxFCl is sent in the packet.

No synchronization is required of the test ID value between initiating and responding MEPs because the test ID is configured at the initiating MEP, and the responding MEP uses the test ID it receives from the initiating MEP. Because ETH-SLM is a sampling technique, it is less precise than counting the service frames. Also, the accuracy of measurement depends on the number of SLM frames used or the period for transmitting SLM frames.

Reception of SLMs and Transmission of SLRs

After the destination MEP receives a valid SLM frame from the source MEP, an SLR frame is generated and transmitted to the requesting or source MEP. The SLR frame is valid if the MEG level and the destination MAC address match the receiving MEP's MAC address. All the fields in the SLM PDUs are copied from the SLM request except for the following fields:

- The source MAC address is copied to the destination MAC address and the source address contains the MEP's MAC address.
- The value of the OpCode field is changed from SLM to SLR (54).
- The responder MEP ID is populated with the MEP's MEP ID.
- TxFCb is saved with the value of the local counter RxFCl at the time of SLR frame transmission.
- An SLR frame is generated every time an SLM frame is received; therefore, RxFCl in the responder is equal to the number of SLM frames received and also equal to the number of SLR frames sent. At the responder or receiving MEP, RxFCl equals TxFCl.

Reception of SLRs

After an SLM frame (with a given TxFCf value) is transmitted, a MEP expects to receive a corresponding SLR frame (carrying the same TxTCf value) within the timeout value from its peer MEP. SLR frames that are received after the timeout value (5 seconds) are discarded. With the information contained in SLR frames, a MEP determines the frame loss for the specified measurement period. The measurement period is a time interval during which the number of SLM frames transmitted is statistically adequate to make a measurement at a given accuracy. A MEP uses the following values to determine near-end and far-end frame loss during the measurement period:

- Last received SLR frame's TxFCf and TxFCb values and the local counter RxFCl value at the end of the measurement period. These values are represented as TxFCf[tc], TxFCb[tc], and RxFCl[tc], where tc is the end time of the measurement period.
- SLR frame's TxFCf and TxFCb values of the first received SLR frame after the test starts and local counter RxFCl at the beginning of the measurement period. These values are represented as TxFCf[tp], TxFCb[tp], and RxFCl[tp], where tp is the start time of the measurement period.

 Copyright © 2013, Juniper Networks, Inc.

For each SLR packet that is received, the local RxFCl counter is incremented at the sending or source MEP.

Computation of Frame Loss

Synthetic frame loss is calculated at the end of the measurement period on the basis of the value of the local counters and the information from the last frame received. The last received frames contains the TxFCf and TxFCb values. The local counter contains the RxFCl value. Using these values, frame loss is determined using the following formula:

Frame loss (far-end) = TxFCf – TxFCb

Frame loss (near-end) = TxFCb – RxFCl

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Format of ETH-SLM Messages on page 317
- Guidelines for Configuring ETH-SLM on page 321
- Scenarios for Configuration of ETH-SLM on page 316
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Starting an On-Demand ETH-SLM Session on page 343
- Troubleshooting Failures with ETH-SLM on page 348
- *Ethernet Interfaces*

Guidelines for Configuring ETH-SLM

Keep the following points in mind when you configure the ETH-SLM functionality:

- The monitoring application for Ethernet OAM is initiated in the master Routing Engine. When a stateful switchover process occurs, the monitoring application is disabled. For on-demand ETH-SLM, graceful Routing Engine switchover (GRES) support is not applicable. For proactive ETH-SLM, the service-level agreement (SLA) iterators are restored during a stateful switchover process. If the adjacencies do not time out, the ETH-SLM statistics are preserved and proactive ETH-SLM supports GRES.
- ETH-SLM is initiated only when the MEP session is up. Unified in-service software upgrade (ISSU) support for ETH-SLM depends on the unified ISSU support for CFM. For CFM, unified ISSU is supported using the loss threshold TLV to avoid CFM connectivity loss during the upgrade. The receiving or the destination MEP increases the threshold time during the termination of sessions. If you start a unified ISSU operation when on-demand ETH-SLM is in progress, the SLM request and reply messages are lost at the local Packet Forwarding Engine.

  When an on-demand ETH-SLM is requested, if the local source MEP undergoes a unified ISSU, a message is displayed stating that the MEP is undergoing a unified ISSU. If the remote MEP is undergoing a unified ISSU (detected through the loss threshold TLV), a message is displayed stating that the remote MEP is undergoing a unified ISSU.

Copyright © 2013, Juniper Networks, Inc.

Also, if it is not possible to identify whether unified ISSU is in progress on a remote MEP, the SLM packets are lost at the system where unified ISSU is in progress and the loss calculation results do not provide a valid cause for the loss. Unified ISSU is not supported for both on-demand and proactive ETH-SLM.

- The maximum number of SLA iterator profiles that can be configured in the system is 255.
- ETH-SLM is not supported for virtual private LAN service (VPLS) (point-to-multipoint measurements are not supported). The ETH-SLM frames are not generated with multicast class 1 destination address. Similarly, ETH-SLM does not respond to ETH-SLM requests with multicast DA. ETH-SLM for VPLS for point-to-point Ethernet connection is supported using directed unicast destination MAC addresses, although point-to-multipoint topologies are not supported.
- A unicast destination address may be used in provisioned environments for point-to-point connections. However, it requires that the unicast destination address of the downstream MEP must have been configured on the MEP transmitting an alarm indication signal (AIS).
- ETH-SLM is not supported on aggregated Ethernet (ae) interfaces and on downstream MEPs on label-switched interfaces (LSIs).
- The number of ETH-SLM sessions for proactive ETH-SLM that can be supported is limited to the total number of iterators that can be supported in the system. This limitation includes the iterator support for other measurement types such as loss, statistical frame loss, and two-way delay. A new iterator type, SLM, is added to support ETH-SLM. The total number of SLA iterators that you can configure in the system is equal to the total number of iterations supported in the system.
- For on-demand SLM, the minimum period between two SLM requests is 100 milliseconds.
- For proactive SLM, the minimum period between two SLM requests is 10 milliseconds for distributed mode and 100 milliseconds for non-distributed mode.
- ETH-SLM frames are always marked as drop-ineligible in compliance with the ITU-T Y.1731 standard.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Transmission of ETH-SLM Messages on page 319
- Format of ETH-SLM Messages on page 317
- Scenarios for Configuration of ETH-SLM on page 316
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Starting an On-Demand ETH-SLM Session on page 343
- Troubleshooting Failures with ETH-SLM on page 348
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Configuring an Iterator Profile

You can create an iterator profile with its parameters to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for delay measurement or loss measurement.

To create an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   [edit]
   user@host# **edit protocols oam ethernet connectivity-fault-management performance-monitoring**

2. Configure the SLA measurement monitoring iterator:

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# **edit sla-iterator-profiles**

3. Configure an iterator profile—for example, i1:

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
   user@host# **set i1**

4. (Optional) Configure the cycle time, which is the amount of time (in milliseconds) between back-to-back transmission of SLA frames for one connection, with values from 10 through 3,600,000. The default value is 1000 ms.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set cycle-time *cycle-time-value***

5. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value), with values from 1 through 2000. The default value is 2000.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set iteration-period *iteration-period-value***

6. Configure the measurement type as loss measurement, statistical frame-loss measurement, or two-way delay measurement.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set measurement-type (loss | statistical-frame-loss | two-way-delay)**

7. (Optional) Configure the calculation weight for delay with values from 1 through 65,535. The default value is 1 (applicable only for two-way delay measurement).

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set calculation-weight delay *delay-value***

8. (Optional) Configure the calculation weight for delay variation with values from 1 through 65,535. The default value is 1 (applicable only for two-way delay measurement).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
  sla-iterator-profiles i1]
user@host# set calculation-weight delay-variation delay-variation-value
```

9. Configure the **disable** statement to stop the iterator (that is, disable the iterator profile).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
  sla-iterator-profiles i1]
user@host# set disable
```

10. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
  sla-iterator-profiles]
user@host# show i1
  cycle-time cycle-time-value;
  iteration-period iteration-period-value;
  measurement-type (loss | two-way-delay);
  calculation-weight {
    delay delay-weight;
    delay-variation delay-variation-weight;
  }
```

Related Documentation

- Proactive Mode on page 312
- Configuring a Remote MEP with an Iterator Profile on page 324
- *Verifying the Configuration of an Iterator Profile*
- Managing Iterator Statistics on page 353
- *Ethernet Interfaces*

Configuring a Remote MEP with an Iterator Profile

You can associate a remote maintenance association end point (MEP) with more than one iterator profile.

To configure a remote MEP with an iterator profile:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit protocols oam ethernet connectivity-fault-management
  maintenance-domain md-name maintenance-association ma-name mep mep-id
```

2. Configure the remote MEP with values from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name mep mep-id]
user@host# set remote-mep remote-mep-id
```

3. Set the iterator profile.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  md-name maintenance-association ma-name mep mep-id remote-mep
  remote-mep-id]
user@host# set sla-iterator-profile profile-name
```

 Copyright © 2013, Juniper Networks, Inc.

4. (Optional) Set the size of the data TLV portion of the Y.1731 data frame with values from 1 through 1400 bytes. The default value is 1.

   ```
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     md-name maintenance-association ma-name mep mep-id remote-mep remote-mep-id
     sla-iterator-profile profile-name]
   user@host# set data-tlv-size size
   ```

5. (Optional) Set the iteration count, which indicates the number of iterations for which this connection should partake in the iterator for acquiring SLA measurements, with values from 1 through 65,535. The default value is 0 (that is, infinite iterations).

   ```
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     md-name maintenance-association ma-name mep mep-id remote-mep remote-mep-id
     sla-iterator-profile profile-name]
   user@host# set iteration-count count-value
   ```

6. (Optional) Set the priority, which is the **vlan-pcp** value that is sent in the Y.1731 data frames, with values from 0 through 7. The default value is 0.

   ```
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     md-name maintenance-association ma-name mep mep-id remote-mep remote-mep-id
     sla-iterator-profile profile-name]
   user@host# set priority priority-value
   ```

7. Verify the configuration.

   ```
   [edit protocols oam ethernet connectivity-fault-management maintenance-domain
     md-name maintenance-association ma-name mep mep-id remote-mep
     remote-mep-id]
   user@host# show
   sla-iterator-profile profile-name {
     data-tlv-size size;
     iteration-count count-value;
     priority priority-value;
   }
   ```

Related Documentation

- Proactive Mode on page 312
- Configuring an Iterator Profile on page 323
- *Verifying the Configuration of an Iterator Profile*
- Managing Iterator Statistics on page 353
- *Ethernet Interfaces*

Configuring Statistical Frame Loss Measurement for VPLS Connections

Using proactive statistical frame loss measurement, you can monitor VPLS connections on MX Series routers . Statistical frame loss measurement allows you to monitor the quality of Ethernet connections for service level agreements (SLAs). Point-to-point and multipoint-to-multipoint connections configured on MX Series routers can be monitored by registering the connection on an iterator and initiating periodic SLA measurement of frame transmissions on the connections.

Copyright © 2013, Juniper Networks, Inc.

Iterators periodically transmit SLA measurement packets using ITU-Y.1731 compliant frames. The iterator sends periodic measurement packets for each of the connections registered to it. These measurement cycles are transmitted in such a way as to not overlap, reducing the processing demands placed on the CPU. The measurement packets are exchanged between the source user network interface (UNI) port and the destination UNI port, providing a sequence of timed performance measurements for each UNI pair. The Frame Loss Ratio (FLR) and connection availability can be computed from these measurements using statistics.

The following steps outline how to configure statistical frame loss measurement for VPLS connections:

1. To configure proactive ETH-DM measurement for a VPLS connection, see "Guidelines for Configuring Routers to Support an ETH-DM Session" on page 326.
2. To enable statistical loss measurement for a VPLS connection, configure an iterator for the VPLS connection using the *sla-iterator-profiles* statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level. For detailed instructions, see "Configuring an Iterator Profile" on page 323.
3. As part of the iterator configuration, include the **statistical-frame-loss** option for the *measurement-type* statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]** hierarchy level.
4. Once you have enabled the iterator, you can display the statistical frame loss for a VPLS connection by issuing the **show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator *identifier* maintenance-domain *name* maintenance-association *name* local-mep *identifier* remote-mep *identifier*** command.

Related Documentation

- Guidelines for Configuring Routers to Support an ETH-DM Session on page 326
- Configuring an Iterator Profile on page 323

- *Verifying the Configuration of an Iterator Profile*
- *Ethernet Interfaces*

Guidelines for Configuring Routers to Support an ETH-DM Session

Keep the following guidelines in mind when configuring routers to support an Ethernet frame delay measurement (ETH-DM) session:

- Configuration Requirements for ETH-DM on page 327
- Configuration Options for ETH-DM on page 327

 Copyright © 2013, Juniper Networks, Inc.

Configuration Requirements for ETH-DM

You can obtain ETH-DM information for a link that meets the following requirements:

- The measurements can be performed between peer maintenance association endpoints (MEPs) on two routers.
- The two MEPs must be configured on two Ethernet physical interfaces or on two Ethernet logical interfaces. For more information, see "Configuring a Maintenance Endpoint" on page 259.
- The two MEPs must be configured—on their respective routers—under the same maintenance association (MA) identifier. For more information, see "Creating a Maintenance Association" on page 255.
- On both routers, the MA must be associated with the same maintenance domain (MD) name. For more information, see "Creating the Maintenance Domain" on page 252.
- On both routers, periodic packet management (PPM) must be running on the Routing Engine and Packet Forwarding Engine, which is the default configuration. You can disable PPM on the Packet Forwarding Engine only. However, the Ethernet frame delay measurement feature requires that distributed PPM remain enabled on the Packet Forwarding Engine of both routers. For more information about **ppm**, see the *Junos OS Routing Protocols Library for Routing Devices*.
- If the PPM process (**ppm**) is disabled on the Packet Forwarding Engine, you must re-enable it. Re-enabling distributed **ppm** entails restarting the **ethernet-connectivity-fault-management** process, which causes all connectivity fault management (CFM) sessions to re-establish. For more information about CFM sessions, see "Configuring Ethernet Local Management Interface" on page 268.

NOTE: The Ethernet frame delay measurement feature is supported only for MEPs configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.

Configuration Options for ETH-DM

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 304
- Configuring Routers to Support an ETH-DM Session on page 334
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Guidelines for Starting an ETH-DM Session

Keep the following guidelines in mind when preparing to start an Ethernet frame delay measurement (ETH-DM) session:

ETH-DM Session Prerequisites

Before you can start an ETH-DM session, you must configure two MX Series routers to support ETH-DM by defining the two CFM-enabled physical or logical Ethernet interfaces on each router. This entails creating and configuring CFM maintenance domains, maintenance associations, and maintenance association end points on each router. For more information about enabling CFM on an Ethernet interface, see "Creating the Maintenance Domain" on page 252.

NOTE: **The Ethernet frame delay measurement feature is supported only for maintenance association end points configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.**

For specific information about configuring routers to support ETH-DM, see "Guidelines for Configuring Routers to Support an ETH-DM Session" on page 326 and "Configuring Routers to Support an ETH-DM Session" on page 334.

ETH-DM Session Parameters

You can initiate a one-way or two-way ETH-DM session by entering the **monitor ethernet delay-measurement** operational command at a router that contains one end of the service for which you want to measure frame delay. The command options specify the ETH-DM session in terms of the CFM elements:

- The type of ETH-DM measurement (one-way or two-way) to be performed.
- The Ethernet service for which the ETH-DM measurement is to be performed:
  - CFM maintenance domain—Name of the existing maintenance domain (MD) for which you want to measure Ethernet frame delays. For more information, see "Creating the Maintenance Domain" on page 252.
  - CFM maintenance association—Name of an existing maintenance association (MA) within the maintenance domain. For more information, see "Creating a Maintenance Association" on page 255.
  - Remote CFM maintenance association end point—The unicast MAC address or the numeric identifier of the remote maintenance association end point (MEP)—the physical or logical interface on the remote router that resides in the specified MD

 Copyright © 2013, Juniper Networks, Inc.

and is named in the specified MA—with which to perform the ETH-DM session. For more information, see "Configuring a Maintenance Endpoint" on page 259.

- Optional specifications:
  - Count—You can specify the number of ETH-DM requests to send for this frame delay measurement session. The range is from 1 through 65,535 frames. The default value is 10 frames.

    NOTE: Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).
  - Frame interval—You can specify the number of seconds to elapse between ETH-DM frame transmittals. The default value is 1 second.

For more detailed information about the parameters you can specify to start an ETH-DM session, see the **monitor ethernet delay-measurement** operational command description in the *Junos OS Operational Mode Commands*.

Restrictions for an ETH-DM Session

The following restrictions apply to an ETH-DM session:

- You cannot run multiple simultaneous ETH-DM sessions with the same remote MEP or MAC address.
- For a given ETH-DM session, you can collect frame delay information for a maximum of 65,535 frames.
- For a given CFM session (pair of peer MEPs), the ETH-DM database stores a maximum of 100 statistics, with the older statistics being "aged out" as newer statistics are collected for that pair of MEPs.
  - For one-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the receiver MEP is defined.
  - For two-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the initiator MEP is defined.

  Depending on the number of frames exchanged in the individual ETH-DM sessions, the ETH-DM database can contain statistics collected through multiple ETH-DM sessions.
- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *Junos OS High Availability Library for Routing Devices*.
- Accuracy of frame delay data is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 304
- Starting an ETH-DM Session on page 338
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 330
- **monitor ethernet delay-measurement** operational command
- *Ethernet Interfaces*

Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts

ETH-DM Statistics

Ethernet frame delay statistics are the frame delay and frame delay variation values determined by the exchange of frames containing ETH-DM protocol data units (PDUs).

- For a one-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, including the exchange of one-way delay PDU frames, see "Ethernet Frame Delay Measurements Overview" on page 304.
- For a two-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, including the exchange of two-way delay PDU frames, see "Ethernet Frame Delay Measurements Overview" on page 304.

A CFM database stores CFM-related statistics and—for Ethernet interfaces that support ETH-DM—the 100 most recently collected ETH-DM statistics for that pair of MEPs. You can view ETH-DM statistics by using the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command to display the CFM statistics for the MEP that collects the ETH-DM statistics you want to view.

Table 25 on page 330 describes the ETH-DM statistics calculated in an ETH-DM session.

Table 25: ETH-DM Statistics

Field Name	Field Description
One-way delay (μsec)†	For a one-way ETH-DM session, the frame delay, in microseconds, collected at the receiver MEP. To display frame delay statistics for a given one-way ETH-DM session, use the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command at the receiver MEP for that session.

 Copyright © 2013, Juniper Networks, Inc.

Table 25: ETH-DM Statistics *(continued)*

Field Name	Field Description
Two-way delay (μsec)	For a two-way ETH-DM session, the frame delay, in microseconds, collected at the initiator MEP. When you start a two-way frame delay measurement, the CLI output displays each DMR frame receipt timestamp and corresponding DMM frame delay and delay variation collected as the session progresses. To display frame delay statistics for a given two-way ETH-DM session, use the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command at the initiator MEP for that session.
Average delay†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Average delay** field displays the average one-way and two- frame delays among all ETH-DM statistics collected at the CFM session level. For example, suppose you start two one-way ETH-DM sessions for 50 counts each, one after the other. If, after both measurement sessions complete, you use a **show** command to display 100 ETH-DM statistics for that CFM session, the **Average delay** field displays the average frame delay among all 100 statistics.
Average delay variation†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay variation among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Average delay variation** field displays the average one-way and two- frame delay variations among all ETH-DM statistics collected at the CFM session level.
Best-case delay†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the lowest two-way frame delay value among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Best case delay** field displays the lowest one-way and two-way frame delays among all ETH-DM statistics collected at the CFM session level.
Worst-case delay†	When you start a two-way frame delay measurement, the CLI output includes a runtime display of the highest two-way frame delay value among the statistics collected for the ETH-DM session only. When you display ETH-DM statistics using a **show** command, the **Worst case delay** field displays the highest one-way and two-way frame delays among all statistics collected at the CFM session level.

†When you start a one-way frame delay measurement, the CLI output displays **NA** (“not available”) for this field. One-way ETH-DM statistics are collected at the remote (receiver) MEP. Statistics for a given one-way ETH-DM session are available only by displaying CFM statistics for the receiver MEP.

Copyright © 2013, Juniper Networks, Inc.

ETH-DM Statistics Retrieval

At the receiver MEP for a one-way session, or at the initiator MEP for a two-way session, you can display all ETH-DM statistics collected at a CFM session level by using the following operational commands:

- **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*> <count *count*>**
- **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*> <count *count*>**

ETH-DM Frame Counts

The number of ETH-DM PDU frames exchanged in a ETH-DM session are stored in the CFM database on each router.

Table 26 on page 332 describes the ETH-DM frame counts collected in an ETH-DM session.

Table 26: ETH-DM Frame Counts

Field Name	Field Description
1DMs sent	Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. Stored in the CFM database of the MEP initiating a one-way frame delay measurement.
Valid 1DMs received	Number of valid 1DM frames received. Stored in the CFM database of the MEP receiving a one-way frame delay measurement.
Invalid 1DMs received	Number of invalid 1DM frames received. Stored in the CFM database of the MEP receiving a one-way frame delay measurement.
DMMs sent	Number of delay measurement message (DMM) PDU frames sent to the peer MEP in this session. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.
DMRs sent	Number of delay measurement reply (DMR) frames sent (in response to a received DMM). Stored in the CFM database of the MEP responding to a two-way frame delay measurement.
Valid DMRs received	Number of valid DMR frames received. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.
Invalid DMRs received	Number of invalid DMR frames received. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.

Copyright © 2013, Juniper Networks, Inc.

ETH-DM Frame Count Retrieval

Each router counts the number of ETH-DM frames sent or received and stores the counts in a CFM database.

Frame Counts Stored in CFM Databases

You can display ETH-DM frame counts for MEPs assigned to specified Ethernet interfaces or for specified MEPs in CFM sessions by using the following operational commands:

- **show oam ethernet connectivity-fault-management interfaces** (**detail** | **extensive**)
- **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*** **<local-mep** ***mep-id*****>** **<remote-mep** ***mep-id*****>**

One-Way ETH-DM Frame Counts

For a one-way ETH-DM session, delay statistics are collected at the receiver MEP only, but frame counts are collected at both MEPs. As indicated in Table 26 on page 332, one-way ETH-DM frame counts are tallied from the perspective of each router in the session:

- At the initiator MEP, the router counts the number of 1DM frames sent.
- At the receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

You can also view one-way ETH-DM frame counts—for a receiver MEP—by using the **show oam ethernet connectivity-fault-management mep-statistics** command to display one-way statistics and frame counts together.

Two-Way ETH-DM Frame Counts

For a two-way ETH-DM session, delay statistics are collected at the initiator MEP only, but frame counts are collected at both MEPs. As indicated in Table 26 on page 332, two-way ETH-DM frame counts are tallied from the perspective of each router in the session:

- At the initiator MEP, the router counts the number of DMM frames sent, valid DMR frames received, and invalid DMR frames received.
- At the responder MEP, the router counts the number of DMR frames sent.

You can also view two-way ETH-DM frame counts—for an initiator MEP—by using the **show oam ethernet connectivity-fault-management mep-statistics** command to display two-way statistics and frame counts together.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 304
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 349
- Example: One-Way Ethernet Frame Delay Measurement on page 358
- **clear oam ethernet connectivity-fault-management statistics** command

- **show oam ethernet connectivity-fault-management mep-statistics** command
- **show oam ethernet connectivity-fault-management delay-statistics** command
- ***show oam ethernet connectivity-fault-management interfaces*** (**detail** | **extensive**) command
- **show oam ethernet connectivity-fault-management mep-database** command
- *Ethernet Interfaces*

Configuring Routers to Support an ETH-DM Session

Configuring MEP Interfaces

Before you can start an Ethernet frame delay measurement session across an Ethernet service, you must configure two MX Series routers to support ETH-DM.

To configure an Ethernet interface on a MX Series router to support ETH-DM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
  ethernet-interface-name {
    vlan-tagging;
    unit logical-unit-number {
      vlan-id vlan-id; # Both interfaces on this VLAN
    }
  }
}
```

   Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md-name { # On both routers
        level number;
        maintenance-association ma-name { # On both routers
          continuity-check {
            interval 100ms;
            hold-interval 1;
          }
          mep mep-id { # Attach to VLAN interface
```

 Copyright © 2013, Juniper Networks, Inc.

```
                    auto-discovery;
                    direction (up | down);
                    interface interface-name;
                    priority number;
                  }
                }
              }
            }
          }
        }
```

Ensuring That Distributed ppm Is Not Disabled

By default, the router's period packet management process (**ppm**) runs sessions distributed to the Packet Forwarding Engine in addition to the Routing Engine. This process is responsible for periodic transmission of packets on behalf of its various client processes, such as Bidirectional Forwarding Detection (BFD), and it also receives packets on behalf of client processes.

In addition, **ppm** handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With **ppm** processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run such processes as BFD on the Packet Forwarding Engine.

Distributed ppm Required for ETH-DM

Ethernet frame delay measurement requires that **ppm** remains distributed to the Packet Forwarding Engine. If **ppm** is not distributed to the Packet Forwarding Engines of both routers, ETH-DM PDU frame timestamps and ETH-DM statistics are not valid.

Before you start ETH-DM, you must verify that the following configuration statement is *NOT* present:

```
[edit]
routing-options {
  ppm {
    no-delegate-processing;
  }
}
```

If distributed **ppm** processing is disabled (as shown in the stanza above) on either router, you must re-enable it in order to use the ETH-DM feature.

Procedure to Ensure that Distributed ppm is Not Disabled

To ensure that distributed **ppm** is not disabled on a router:

1. Display the packet processing management (PPM) configuration to determine whether distributed **ppm** is disabled.

   - In the following example, distributed **ppm** is enabled on the router. In this case, you do not need to modify the router configuration:

     ```
     [edit]
     user@host# show routing-options
     ppm;
     ```

- In the following example, distributed **ppm** is disabled on the router. In this case, you must proceed to Step 2 to modify the router configuration:

```
[edit]
user@host show routing-options
ppm {
   no-delegate-processing;
}
```

2. Modify the router configuration to re-enable distributed **ppm** and restart the Ethernet OAM Connectivity Fault Management process *ONLY IF* distributed **ppm** is disabled (as determined in the previous step).

   a. Before continuing, make any necessary preparations for the possible loss of connectivity on the router.

      Restarting the **ethernet-connectivity-fault-management** process has the following effect on your network:

      - All connectivity fault management (CFM) sessions re-establish.
      - All ETH-DM requests on the router terminate.
      - All ETH-DM statistics and frame counts reset to 0.

   b. Modify the router configuration to re-enable distributed **ppm**. For example:

```
[edit]
user@host# delete routing-options ppm no-delegate-processing
```

   c. Commit the updated router configuration. For example:

```
[edit]
user@host# commit and-quit
commit complete
exiting configuration mode
```

   d. To restart the Ethernet OAM Connectivity-Fault-Management process, enter the **restart ethernet-connectivity-fault-management <gracefully | immediately | soft>** operational mode command. For example:

```
user@host> restart ethernet-connectivity-fault-management
Connectivity fault management process started, pid 9893
```

Enabling the Hardware-Assisted Timestamping Option

By default, Ethernet frame delay measurement uses software for timestamping transmitted and received ETH-DM frames. For Ethernet interfaces, you can optionally use hardware timing to assist in the timestamping of received ETH-DM frames to increase the accuracy of delay measurements.

Enabling hardware-assisted timestamping of received frames can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction.

 Copyright © 2013, Juniper Networks, Inc.

To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      performance-monitoring {
        hardware-assisted-timestamping;
      }
    }
  }
}
```

Configuring the Server-Side Processing Option

You can delegate the server-side processing (for both two-way delay measurement and loss measurement) to the Packet Forwarding Engine to prevent overloading on the Routing Engine. By default, the server-side processing is done by the Routing Engine.

To configure the server-side processing option:

1. In configuration mode, go to the following hierarchy level:

   ```
   user@host# edit protocols oam ethernet connectivity-fault-management
     performance-monitoring
   ```

2. Configure the server-side processing option.

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# set delegate-server-processing
   ```

3. Verify the configuration.

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@host# show
   performance-monitoring {
     delegate-server-processing;
   }
   ```

Related Documentation

- On-Demand Mode on page 311
- *Ethernet Interfaces*
- Ethernet Frame Delay Measurements Overview on page 304
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 326
- *Ethernet Interfaces*
- Ethernet Frame Delay Measurements Overview on page 304
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 326
- *Ethernet Interfaces*

Starting an ETH-DM Session

Using the monitor ethernet delay-measurement Command

After you have configured two MX Series routers to support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), you can initiate a one-way or two-way Ethernet frame delay measurement session from the CFM maintenance association end point (MEP) on one of the routers to the peer MEP on the other router.

To start an ETH-DM session between the specified local MEP and the specified remote MEP, enter the **monitor ethernet delay-measurement** command at operational mode. The syntax of the command is as follows:

monitor ethernet delay-measurement
(one-way | two-way)
maintenance-domain *md-name*
maintenance-association *ma-name*
(*remote-mac-address* | mep *remote-mep-id*)
<count *frame-count*>
<wait *interval-seconds*>
<priority *802.1p value*>
<size>
<no-session-id-tlv>
<xml>

For a one-way frame delay measurement, the command displays a runtime display of the number of 1DM frames sent from the initiator MEP during that ETH-DM session. One-way frame delay and frame delay variation measurements from an ETH-DM session are collected in a CFM database at the router that contains the receiver MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

For a two-way frame delay measurement, the command displays two-way frame delay and frame delay variation values for each round-trip frame exchange during that ETH-DM session, as well as a runtime display of useful summary information about the session: average delay, average delay variation, best-case delay, and worst-case delay. Two-way frame delay and frame delay variation values measurements from an ETH-DM session are collected in a CFM database at the router that contains the initiator MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

NOTE: **Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).**

For a complete description of the **monitor ethernet delay-measurement** operational command, see the *Junos OS Operational Mode Commands*.

 Copyright © 2013, Juniper Networks, Inc.

Starting a One-Way ETH-DM Session

To start a one-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement one-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

Starting a Two-Way ETH-DM Session

To start a two-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement two-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec  Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Ethernet Frame Delay Measurements Overview on page 304
- Guidelines for Starting an ETH-DM Session on page 328
- **monitor ethernet delay-measurement** command
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 330
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 349
- *Ethernet Interfaces*

Starting a Proactive ETH-SLM Session

To start a proactive Ethernet synthetic loss measurement (ETH-SLM) session, you must configure the Ethernet interfaces on maintenance association end points (MEPs) on which packets transmitted with synthetic frame loss need to be analyzed. You must then create an iterator profile to transmit service-level agreement (SLA) measurement packets for ETH-SLM and associate the local and remote MEPs with the profile.

Configuring MEP Interfaces

Before you can start an Ethernet synthetic frame loss measurement session across an Ethernet service, you must configure two ACX Series routers to support ETH-SLM.

To configure an Ethernet interface on an ACX Series router to support ETH-SLM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
  ethernet-interface-name {
    vlan-tagging;
    unit logical-unit-number {
      vlan-id vlan-id; # Both interfaces on this VLAN
    }
  }
}
```

   Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md-name { # On both routers
        level number;
```

 Copyright © 2013, Juniper Networks, Inc.

```
                maintenance-association ma-name { # On both routers
                    continuity-check {
                        interval 100ms;
                        hold-interval 1;
                    }
                    mep mep-id { # Attach to VLAN interface
                        auto-discovery;
                        direction (up | down);
                        interface interface-name;
                        priority number;
                    }
                }
            }
        }
    }
}
```

Configuring an Iterator Profile for ETH-SLM

You can create an iterator profile with its parameters to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for synthetic loss measurement.

To create an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   ```
   [edit]
   user@host# edit protocols oam ethernet connectivity-fault-management
     performance-monitoring
   ```

2. Configure the SLA measurement monitoring iterator:

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring]
   user@host# edit sla-iterator-profiles
   ```

3. Configure an iterator profile—for example, i1:

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring
     sla-iterator-profiles]
   user@host# set i1
   ```

4. (Optional) Configure the cycle time, which is the amount of time (in milliseconds) between back-to-back transmission of SLA frames for one connection, with a value from 10 through 3,600,000. The default value is 1000 ms.

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring
     sla-iterator-profiles i1]
   user@host# set cycle-time cycle-time-value
   ```

5. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value), with a value from 1 through 2000. The default value is 2000.

   ```
   [edit protocols oam ethernet connectivity-fault-management performance-monitoring
     sla-iterator-profiles i1]
   user@host# set iteration-period iteration-period-value
   ```

Copyright © 2013, Juniper Networks, Inc.

6. Configure the measurement type as synthetic loss measurement.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set measurement-type slm**

7. Configure the **disable** statement to stop the iterator (that is, disable the iterator profile).

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
   user@host# **set disable**

8. Verify the configuration.

   [edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
   user@host# **show i1**
   cycle-time *cycle-time-value*;
   iteration-period *iteration-period-value*;
   measurement-type slm;

Associating the Iterator Profile with MEPs for ETH-SLM

You can associate a remote maintenance association end point (MEP) with more than one iterator profile.

To configure a remote MEP with an iterator profile:

1. In configuration mode, go to the following hierarchy level:

   user@host# **edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id***

2. Configure the remote MEP ID with a value from 1 through 8191.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]
   user@host# **set remote-mep *remote-mep-id***

3. Set the iterator profile.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]
   user@host# **set sla-iterator-profile *profile-name***

4. (Optional) Set the size of the data TLV portion of the Y.1731 data frame with a value from 1 through 1400 bytes. The default value is 1.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set data-tlv-size *size***

5. (Optional) Set the iteration count, which indicates the number of iterations for which this connection should partake in the iterator for acquiring SLA measurements, with a value from 1 through 65,535. The default value is 0 (that is, infinite iterations).

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]

 Copyright © 2013, Juniper Networks, Inc.

user@host# **set iteration-count** ***count-value***

6. (Optional) Set the priority, which is the **vlan-pcp** value that is sent in the Y.1731 data frames, with a value from 0 through 7. The default value is 0.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id* sla-iterator-profile *profile-name*]
   user@host# **set priority** ***priority-value***

7. Verify the configuration.

   [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]
   user@host# **show**
   sla-iterator-profile *profile-name* {
     data-tlv-size *size*;
     iteration-count *count-value*;
     priority *priority-value*;
   }

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Transmission of ETH-SLM Messages on page 319
- Format of ETH-SLM Messages on page 317
- Guidelines for Configuring ETH-SLM on page 321
- Scenarios for Configuration of ETH-SLM on page 316
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting an On-Demand ETH-SLM Session on page 343
- Troubleshooting Failures with ETH-SLM on page 348

- *Ethernet Interfaces*

Starting an On-Demand ETH-SLM Session

To start an on-demand Ethernet synthetic loss measurement (ETH-SLM) session, type the **monitor ethernet synthetic-loss-measurement one-way** command in operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0
      Synthetic Loss measurement statistics:
       SLM packets sent                          : 100
       SLR packets received                      : 100
       Accumulated SLM statistics:
       Local TXFC1 value                         : 100
       Local RXFC1 value                         : 100
       Last Received SLR frame TXFCf(tc)         : 100
       Last Received SLR frame TXFCb(tc)         : 100
```

```
SLM Frame Loss:
Frame Loss (far-end)                          : 0 (0.00 %)
Frame Loss (near-end)                         : 0 (0.00 %)
```

NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must press **Ctrl + C** to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Transmission of ETH-SLM Messages on page 319
- Format of ETH-SLM Messages on page 317
- Guidelines for Configuring ETH-SLM on page 321
- Scenarios for Configuration of ETH-SLM on page 316
- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Troubleshooting Failures with ETH-SLM on page 348
- *Ethernet Interfaces*

Managing ETH-SLM Statistics and ETH-SLM Frame Counts

Displaying ETH-SLM Statistics Only

Purpose Display on-demand ETH-SLM statistics.

By default, the **show oam ethernet connectivity-fault-management synthetic-loss-statistics** command displays on-demand ETH-SLM statistics for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

Action

- To display the on-demand ETH-SLM statistics collected for MEPs belonging to maintenance association **ma1** within maintenance domain **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md1 maintenance-association ma1**

Copyright © 2013, Juniper Networks, Inc.

- To display the on-demand ETH-SLM statistics collected for ETH-SLM sessions for the local MEP **201** belonging to maintenance association **ma2** within maintenance domain **md2**:

  `user@host>` **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the on-demand ETH-SLM statistics collected for ETH-SLM sessions from local MEPs belonging to maintenance association **ma3** within maintenance domain **md3** to the remote MEP **302**:

  `user@host>` **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md3 maintenance-association ma3 remote-mep 302**

Meaning The output displays on-demand ETH-SLM statistics for MEPs in the specified maintenance association within the specified maintenance domain. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management synthetic-loss-statistics**.

Displaying ETH-SLM Statistics and Frame Counts

Purpose Display on-demand ETH-SLM statistics and ETH-SLM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays on-demand ETH-SLM statistics and frame counts for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

Action

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for MEPs in maintenance association **ma1** within maintenance domain **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1**

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for the local MEP **201** in maintenance association **ma2** within maintenance domain **md2**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for the local MEP in maintenance association **ma3** within maintenance domain **md3** that participates in an ETH-SLM session with the remote MEP **302**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Meaning The output displays on-demand ETH-SLM statistics and ETH-SLM frame counts for MEPs in the specified maintenance association within the specified maintenance domain. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management mep-statistics**.

Displaying ETH-SLM Frame Counts for MEPs by Enclosing CFM Entity

Purpose Display on-demand ETH-SLM frame counts for CFM maintenance association end points (MEPs).

Copyright © 2013, Juniper Networks, Inc.

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

NOTE: **At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display the ETH-SLM frame counts and not the MEP database details.**

Action

- To display CFM database information (including ETH-SLM frame counts) for all MEPs in MA **ma1** within maintenance domain **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma1 maintenance-association ma1**

- To display CFM database information (including ETH-SLM frame counts) only for the local MEP **201** in MA **ma1** within maintenance domain **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display CFM database information (including ETH-SLM frame counts) only for the remote MEP **302** in MA **ma3** within maintenance domain **md3**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Meaning

The output displays ETH-SLM frame counts for MEPs within a particular maintenance domain, or for a specific local or remote MEP. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management mep-database**.

Displaying ETH-SLM Frame Counts for MEPs by Interface or Domain Level

Purpose

Display on-demand ETH-SLM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-SLM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

NOTE: **At the router attached to the initiator MEP, you can only display the ETH-SLM frame counts and not the MEP database details.**

Action

- To display CFM database information (including ETH-SLM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces detail**

 Copyright © 2013, Juniper Networks, Inc.

- To display CFM database information (including ETH-SLM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail**

- To display CFM database information (including ETH-SLM frame counts) only for MEPs enclosed within CFM maintenance domains at level **6**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces level 6 detail**

Meaning The output displays ETH-SLM frame counts for MEPs for the specified interface. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management interfaces**.

Clearing ETH-SLM Statistics and Frame Counts

Purpose Clear the on-demand ETH-SLM statistics and ETH-SLM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the on-demand ETH-SLM statistics and ETH-SLM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management synthetic-loss-measurement**

- To clear the on-demand ETH-SLM statistics and ETH-SLM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

  `user@host>` **clear oam ethernet connectivity-fault-management synthetic-loss-measurement ge-0/5/9.0**

Clearing Iterator Statistics

Purpose Clear the existing iterator statistics and proactive ETH-SLM counters.

Multiple iterators can be associated with remote MEP. However, by default, only one result pertaining to one iterator profile can be cleared.

Action

- To clear the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep *1***

- To clear the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1**

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- **clear oam ethernet connectivity-fault-management synthetic-loss-measurement**
- **show oam ethernet connectivity-fault-management synthetic-loss-statistics**
- **show oam ethernet connectivity-fault-management interfaces** (**detail** | **extensive**)
- **show oam ethernet connectivity-fault-management mep-statistics**
- **show oam ethernet connectivity-fault-management mep-database**
- *Ethernet Interfaces*

Troubleshooting Failures with ETH-SLM

Problem The Ethernet synthetic loss measurement (ETH-SLM) application is not working properly for calculation of frame loss using synthetic frames instead of data traffic

Solution Perform the following steps to analyze and debug any problems with the ETH-SLM functionality.

1. Ensure that ETH-SLM is configured (either proactive or on-demand) to initiate SLM frames. Verify the configuration settings.
2. Examine any failures that might have occurred in the CFM session for which the ETH-SLM feature is enabled. The CFM session must be in the up state for the ETH-SLM functionality to work correctly. Use the **show oam ethernet connectivity-fault-management mep-database maintenance-domain** ***md-name*** **maintenance-association** ***ma-name*** **local-mep** ***mep-id*** **remote-mep** ***remote-mep-id*** command to verify whether the CFM session is in the up state.
3. If the MEP sessions are active, use the appropriate show command to verify the ETH-SLM statistics and to analyze if ETH-SLM frames are transmitted or received.
4. If the transmission of ETH-SLM frames does not happen correctly after you attempt all of the preceding troubleshooting steps, enable the tracing operations for Ethernet CFM by including the **traceoptions** statement at the [**edit protocols oam ethernet connectivity-fault-management**] hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
traceoptions {
  file <filename> <files number  <match regular-expression microsecond-stamp>>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Related Documentation

- Ethernet Synthetic Loss Measurement Overview on page 315
- Transmission of ETH-SLM Messages on page 319
- Format of ETH-SLM Messages on page 317
- Guidelines for Configuring ETH-SLM on page 321
- Scenarios for Configuration of ETH-SLM on page 316

 Copyright © 2013, Juniper Networks, Inc.

- Managing ETH-SLM Statistics and ETH-SLM Frame Counts on page 344
- Starting a Proactive ETH-SLM Session on page 340
- Starting an On-Demand ETH-SLM Session on page 343
- *Ethernet Interfaces*

Managing ETH-DM Statistics and ETH-DM Frame Counts

Displaying ETH-DM Statistics Only

Purpose Display ETH-DM statistics.

By default, the **show oam ethernet connectivity-fault-management delay-statistics** command displays ETH-DM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain ma1 maintenance-association ma1**

- To display the ETH-DM statistics collected for ETH-DM sessions for the local MEP **201** belonging to MA **ma2** and within MD **md2**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the ETH-DM statistics collected for ETH-DM sessions from local MEPs belonging to MA **ma3** and within MD **md3** to remote MEP **302**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md3 maintenance-association ma3 remote-mep 302**

Displaying ETH-DM Statistics and Frame Counts

Purpose Display ETH-DM statistics and ETH-DM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays ETH-DM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1**

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP **201** in MA **ma2** and within MD **md2**:

  user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP in MD **md3** and within MA **ma3** that participates in an ETH-DM session with the remote MEP **302**:

  user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

  user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma1 maintenance-association ma1**

- To display CFM database information (including ETH-DM frame counts) only for local MEP **201** in MA **ma1** within MD **md1**:

  user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain md2 maintenance-association ma2 local-mep 201**

- To display CFM database information (including ETH-DM frame counts) only for remote MEP **302** in MD **md3** within MA **ma3**:

  user@host> **show oam ethernet connectivity-fault-management mep-database maintenance-domain ma3 maintenance-association ma3 remote-mep 302**

Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-DM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

 Copyright © 2013, Juniper Networks, Inc.

NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces detail**
- To display CFM database information (including ETH-DM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail**
- To display CFM database information (including ETH-DM frame counts) only for MEPs enclosed within CFM maintenance domains (MDs) at level **6**:

  `user@host>` **show oam ethernet connectivity-fault-management interfaces level 6 detail**

Clearing ETH-DM Statistics and Frame Counts

Purpose Clear the ETH-DM statistics and ETH-DM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the ETH-DM statistics and ETH-DM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management statistics**
- To clear the ETH-DM statistics and ETH-DM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

  `user@host>` **clear oam ethernet connectivity-fault-management statistics ge-0/5/9.0**

Related Documentation

- *clear oam ethernet connectivity-fault-management statistics*
- *show oam ethernet connectivity-fault-management delay-statistics*
- *show oam ethernet connectivity-fault-management interfaces*
- *show oam ethernet connectivity-fault-management mep-statistics*
- *show oam ethernet connectivity-fault-management mep-database*
- *Ethernet Interfaces*

Managing ETH-LM Statistics

Copyright © 2013, Juniper Networks, Inc.

Displaying ETH-LM Statistics

Purpose Display the ETH-LM statistics.

By default, the **show oam ethernet connectivity-fault-management loss-statistics maintenance-domain *md-name* maintenance-association *ma-name*** command displays ETH-LM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

The following list consists of the CFM-related operational mode commands that have been enhanced to display ETH-LM statistics:

- The **show oam ethernet connectivity-fault-management interfaces detail** command is enhanced to display ETH-DM and ETH-LM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).
- The **show oam ethernet connectivity-fault-management mep-statistics** command is enhanced to display ETH-DM and ETH-LM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).
- The **show oam ethernet connectivity-fault-management mep-database** command is enhanced to display ETH-DM and ETH-LM frame counters for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-LM statistics for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **show oam ethernet connectivity-fault-management loss-statistics**
- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1**
- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1**
- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management mep-database maintenance-domain md1 maintenance-association ma1**

Clearing ETH-LM Statistics

Purpose Clear the ETH-LM statistics.

By default, statistics are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

 Copyright © 2013, Juniper Networks, Inc.

Action

- To clear the ETH-LM statistics for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management loss-statistics**

Related Documentation

- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 349

Managing Iterator Statistics

Displaying Iterator Statistics

Purpose

Retrieve and display iterator statistics.

Muliple iterators can be associated with a remote MEP. However, by default, only one result pertaining to one iterator profile is displayed.

Action

- To display the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1** (here, the iterator profile **i1** is configured for two-way delay measurement):

  `user@host>` **show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1**

```
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent                                        : 23898520
DMM skipped for threshold hit                   : 11000
DMM skipped for threshold hit window            : 0
DMR received                                    : 23851165
DMR out of sequence                             : 1142
DMR received with invalid time stamps           : 36540
Average two-way delay                           : 129 usec
Average two-way delay variation                 : 15 usec
Average one-way forward delay variation         : 22 usec
Average one-way backward delay variation        : 22 usec
Weighted average two-way delay                  : 134 usec
Weighted average two-way delay variation        : 8 usec
Weighted average one-way forward delay variation  : 6 usec
Weighted average one-way backward delay variation : 2 usec
```

  Output fields are listed in the approximate order in which they appear.

Table 27: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields

Output Field Name	Output Field Description
Maintenance domain	Maintenance domain name.
Level	Maintenance domain level configured.
Maintenance association	Maintenance association name.
Local MEP id	Numeric identifier of the local MEP.
Remote MEP id	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
Delay weight	Calculation weight of delay.
Delay variation weight	Calculation weight of delay variation.
DMM sent	Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.
DMM skipped for threshold hit	Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.
DMM skipped for threshold hit window	Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
DMR received	Number of delay measurement reply (DMR) frames received.
DMR out of sequence	Total number of DMR out of sequence packets received.
DMR received with invalid time stamps	Total number of DMR frames received with invalid timestamps.

 Copyright © 2013, Juniper Networks, Inc.

Table 27: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields *(continued)*

Output Field Name	Output Field Description
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Average one-way forward delay variation	Average one-way forward delay variation for the statistics displayed in microseconds.
Average one-way backward delay variation	Average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average two-way delay	Weighted average two-way delay for the statistics displayed in microseconds.
Weighted average two-way delay variation	Weighted average two-way delay variation for the statistics displayed in microseconds.
Weighted average one-way forward delay variation	Weighted average one-way forward delay variation for the statistics displayed in microseconds.
Weighted average one-way backward delay variation	Weighted average one-way backward delay variation for the statistics displayed in microseconds.

- To display the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1** (here, the iterator profile **i1** is configured for loss measurement):

```
user@host> show oam ethernet connectivity-fault-management sla-iterator-statistics
sla-iterator i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1
remote-mep 1
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i2, Iterator Id: 2
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:25 ago)
Reset reason: Adjacency flap

Iterator loss measurement statistics:
LMM sent                               : 238970
LMM skipped for threshold hit          : 60
LMM skipped for threshold hit window   : 0
LMR received                           : 238766
LMR out of sequence                    : 43

Accumulated transmit statistics:
Near-end (CIR)                         : 0
Far-end (CIR)                          : 0
Near-end (EIR)                         : 0
Far-end (EIR)                          : 0
```

```
Accumulated loss statistics:
Near-end (CIR)                       : 0 (0.00%)
Far-end (CIR)                        : 0 (0.00%)
Near-end (EIR)                       : 0 (0.00%)
Far-end (EIR)                        : 0 (0.00%)

Last loss measurement statistics:
Near-end (CIR)                       : 0
Far-end (CIR)                        : 0
Near-end (EIR)                       : 0
Far-end (EIR)                        : 0
```

Output fields are listed in the approximate order in which they appear.

Table 28: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields

Output Field Name	Output Field Description
Maintenance domain	Maintenance domain name.
Level	Maintenance domain level configured.
Maintenance association	Maintenance association name.
Local MEP id	Numeric identifier of the local MEP.
RemoteMEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
LMM sent	Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.
LMM skipped for threshold hit	Number of LMM frames sent to the peer MEP in this session skipped during threshold hit.

Copyright © 2013, Juniper Networks, Inc.

Table 28: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields *(continued)*

Output Field Name	Output Field Description
LMM skipped for threshold hit window	Number of LMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
LMR received	Number of LMRs frames received.
LMR out of sequence	Total number of LMR out of sequence packets received.
Near-end (CIR)	Frame loss associated with ingress data frames for the statistics displayed.
Far-end (CIR)	Frame loss associated with egress data frames for the statistics displayed.
Near-end (EIR)	Frame loss associated with ingress data frames for the statistics displayed.
Far-end (EIR)	Frame loss associated with egress data frames for the statistics displayed.

Clearing Iterator Statistics

Purpose

Clear iterator statistics.

Multiple iterators can be associated with remote MEP. However, by default, only one result pertaining to one iterator profile can be cleared.

Action

- To clear the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep *1***

- To clear the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1**:

  `user@host>` **clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1**

Related Documentation

- Configuring an Iterator Profile on page 323
- Configuring a Remote MEP with an Iterator Profile on page 324
- *Verifying the Configuration of an Iterator Profile*
- Proactive Mode on page 312

Copyright © 2013, Juniper Networks, Inc.

Managing Continuity Measurement Statistics

Displaying Continuity Measurement Statistics

Purpose

Display continuity measurement.

The **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1** command is enhanced to display continuity measurement statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

  `user@host>` **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1**

Clearing Continuity Measurement Statistics

Purpose

Clear the continuity measurement statistics

By default, statistics are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

Action

- To clear the continuity measurement statistics for all MEPs attached to CFM-enabled interfaces on the router:

  `user@host>` **clear oam ethernet connectivity-fault-management continuity-measurement maintenance-domain *md-name* maintenance-association *ma-name* local-mep *local-mep-id* remote-mep *remote-mep-id***

Related Documentation

- *clear oam ethernet connectivity-fault-management continuity-measurement*
- *Ethernet Interfaces*
- *show oam ethernet connectivity-fault-management delay-statistics*

Example: One-Way Ethernet Frame Delay Measurement

Description of the One-Way Frame Delay Measurement Example

This example shows how you can configure two MX Series routers (**MX-PE1** and **MX-PE2**) to support an ETH-DM session between two peer MEPs (MEP **201** and MEP **101**), initiate a one-way ETH-DM session (from MEP **101** to MEP **201**), and then display the ETH-DM

 Copyright © 2013, Juniper Networks, Inc.

statistics and frame counts collected. To increase the accuracy of the ETH-DM statistics, enable optional hardware-assisted timestamping of received ETH-DM frames on the router that contains the receiver MEP.

Routers Used in This Example

To support one-way ETH-DM with optional hardware timestamping of frames on the reception path, the routers used in this example are configured as follows:

- Routers **MX-PE1** and **MX-PE2** are MX Series routers.
- The system clocks of routers **MX-PE1** and **MX-PE2** are closely synchronized.
- On router **MX-PE1**, interface **ge-5/2/9** is an Ethernet port. The traffic load received is heavy.
- On router **MX-PE2**, interface **ge-0/2/5** is an Ethernet port.

ETH-DM Frame Counts for this Example

Both routers count the number of ETH-DM frames sent and received by the peer MEPs in the session and store the frame counts in the CFM databases as follows:

- At router **MX-PE2**, which contains the initiator MEP **101**, the CFM database stores the ETH-DM frame counts for a one-way ETH-DM initiator (the count of 1DM frames sent).
- At router **MX-PE1**, which contains the receiver MEP **201**, the CFM database stores the ETH-DM frame counts for a one-way ETH-DM receiver (the count of valid 1DM frames received and the count of invalid 1DM frames received).

ETH-DM Statistics for this Example

For a one-way frame delay measurement, only the router that contains the receiver MEP measures and stores frame delay statistics. In this example, ETH-DM statistics collected for the session are available only at router **MX-PE1**.

Copyright © 2013, Juniper Networks, Inc.

Steps for the One-Way Frame Delay Measurement Example

The following steps describe an example one-way Ethernet frame delay measurement:

1. At router **MX-PE1**, configure MEP **201** as a CFM maintenance association endpoint in CFM maintenance domain **md6** as follows:

   a. Define the maintenance domain **md6** by associating it with maintenance domain level **6** and maintenance association identifier **ma6**.

   b. Configure the maintenance association by specifying continuity protocol options and specifying MEP identifier **201**.

   c. Configure MEP **201** by attaching it to logical interface **ge-5/2/9.0**, which is a single-tag interface on VLAN **512**.

   The following configuration is only a partial example of a complete and functional router configuration:

```
[edit]
interfaces { # Configure a single-tag logical interface on VLAN 512
  ge-5/2/9 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        maintenance-domain md6 { # Define MD 'md6' on router MX-PE1
          level 6;
          maintenance-association ma6 { # Configure MA 'ma6' on router MX-PE1
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 201 { # Configure MEP 201 on router MX-PE1
              interface ge-5/2/9.0; # Attach to logical interface on VLAN 512
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

 Copyright © 2013, Juniper Networks, Inc.

2. At router **MX-PE2**, configure MEP **101** as a CFM maintenance association endpoint in CFM maintenance domain **md6** as follows:

   a. Define the maintenance domain **md6** by associating it with maintenance domain level **6** and maintenance association identifier **ma6**.

   b. Configure the maintenance association by specifying continuity protocol options and specifying MEP identifier **101**.

   c. Configure MEP **101** by attaching it to logical interface **ge-0/2/5.0**, which is a single-tag interface on VLAN **512**.

   The following configuration is only a partial example of a complete and functional configuration for router **MX-PE2**:

```
[edit]
interfaces { # Configure a single-tag logical interface on VLAN 512
  ge-0/2/5 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        maintenance-domain md6 { # Define MD 'md6' on router MX-PE2
          level 6;
          maintenance-association ma6 { # Configure MA 'ma6' on router MX-PE2
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 101 { # Configure MEP 101 on router MX-PE2
              interface ge–0/2/5.0; # Attach to logical interface on VLAN 512
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
```

3. (Optional) To increase the accuracy of the ETH-DM statistics, modify the configuration of router **MX-PE1**, which contains the receiver MEP, by enabling hardware-assisted timestamping of **1DM** frames received on the router.

```
[edit protocols]
oam {
```

```
      ethernet {
        connectivity-fault-management {
          performance-monitoring {
            hardware-assisted-timestamping;
          }
        }
      }
    }
```

4. At router **MX-PE2**, start a one-way frame delay measurement session from local MEP **101** to remote MEP **201** on router **MX-PE1**:

```
user@MX-PE2> monitor ethernet delay-measurement one-way mep 201 maintenance-domain
md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:90:69:0a:43:94, Interface ge-0/2/5.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

5. At router **MX-PE2**, which contains the initiator MEP, only the ETH-DM frame counts are available. Furthermore, the only frame count tallied for the initiator of a one-way frame delay measurement is the count of 1DM frames transmitted.

   ETH-DM frame counts (the number of 1DM, DMM, and DMR frames exchanged during an ETH-DM session) are stored in the CFM database of both the initiator and receiver MEPs. When you display CFM database information, you can also display the ETH-DM frame counts. You can display CFM database information for all interfaces on the router, or you can limit the output to MEPs associated with certain CFM MDs and MAs.

   - To display CFM database information for MEPs specified by enclosing CFM entities, use the **mep-database** form of the **show oam ethernet connectivity-fault-management** command. A CFM database also stores any ETH-DM frame counts.

     In the example configuration for router **MX-PE2**, MEP **101** is the only MEP defined in MA **ma6** within MD **md6**. Therefore, the **show oam ethernet connectivity-fault management mep-database** command output displays CFM database information for MEP **101** only, even though you do not filter the command output by including the **local-mep** or **remote-mep** command options.

```
user@MX-PE2> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-association ma6
Maintenance domain name: md6, Format: string, Level: 6
  Maintenance association name: ma6, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames

  MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                 : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                         : no
  Statistics:
    CCMs sent                                    : 1590
    CCMs received out of sequence                : 0
```

 Copyright © 2013, Juniper Networks, Inc.

```
    LBMs sent                                    : 0
    Valid in-order LBRs received                 : 0
    Valid out-of-order LBRs received             : 0
    LBRs received with corrupted data            : 0
    LBRs sent                                    : 0
    LTMs sent                                    : 0
    LTMs received                                : 0
    LTRs sent                                    : 0
    LTRs received                                : 0
    Sequence number of next LTM request          : 0
    1DMs sent                                    : 10
    Valid 1DMs received                          : 0
    Invalid 1DMs received                        : 0
    DMMs sent                                    : 0
    DMRs sent                                    : 0
    Valid DMRs received                          : 0
    Invalid DMRs received                        : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      201     00:90:69:0a:43:94       ok    ge-0/2/5.0
```

- To display CFM database information for MEPs specified by interface name, use the **interfaces detail** form of the **show oam ethernet connectivity-fault-management** command. A CFM database also stores any ETH-DM frame counts.

  In the example configuration for router **MX-PE2**, MEP **101** is the only MEP assigned to an interface on the router. Therefore, the **show oam ethernet connectivity-fault-management interfaces** (**detail** | **extensive**) command output displays CFM database information for MEP **101** only, even though you do not filter the command output by including the ***ethernet-interface-name*** or **level** ***md-level*** command options.

```
user@MX-PE2> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
  Maintenance domain name: md6, Format: string, Level: 6
    Maintenance association name: ma6, Format: string
    Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3
frames
    MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
    MEP status: running
    Defects:
      Remote MEP not receiving CCM                 : no
      Erroneous CCM received                       : no
      Cross-connect CCM received                   : no
      RDI sent by some MEP                         : no
    Statistics:
      CCMs sent                                    : 1590
      CCMs received out of sequence                : 0
      LBMs sent                                    : 0
      Valid in-order LBRs received                 : 0
      Valid out-of-order LBRs received             : 0
      LBRs received with corrupted data            : 0
      LBRs sent                                    : 0
      LTMs sent                                    : 0
      LTMs received                                : 0
      LTRs sent                                    : 0
      LTRs received                                : 0
      Sequence number of next LTM request          : 0
      1DMs sent                                    : 10
      Valid 1DMs received                          : 0
```

```
    Invalid 1DMs received                            : 0
    DMMs sent                                        : 0
    DMRs sent                                        : 0
    Valid DMRs received                              : 0
    Invalid DMRs received                            : 0
  Remote MEP count: 1
    Identifier    MAC address        State    Interface
      201     00:90:69:0a:43:94       ok    ge-0/2/5.0
```

NOTE: You can use these same commands—**show oam ethernet connectivity-fault-management mep-database** and **show oam ethernet connectivity-fault-management interfaces (detail | extensive)**—at router **MX-PE1** to display the CFM database information (which includes any ETH-DM frame counts) for receiver MEP 201.

6. At router **MX-PE1**, which contains the receiver MEP, you can use two different **show oam ethernet connectivity-fault-management** commands to display ETH-DM statistics and ETH-DM frame counts.

   - To display only the delay statistics, use the **delay-statistics** form of the **show oam ethernet connectivity-fault-management** command:

     user@MX-PE1> **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md6**

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1

  Remote MAC address: 00:90:69:0a:48:57
    Delay measurement statistics:
    Index  One-way delay  Two-way delay
              (usec)         (usec)
      1        370
      2        357
      3        344
      4        332
      5        319
      6        306
      7        294
      8        281
      9        269
      10       255
    Average one-way delay          : 312 usec
    Average one-way delay variation: 11 usec
    Best case one-way delay        : 255 usec
    Worst case one-way delay       : 370 usec
```

   - To display both the ETH-DM statistics and the CFM database information (which includes any ETH-DM frame counts), use the **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command:

     user@MX-PE1> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md6**

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
  CCMs sent                                    : 3240
  CCMs received out of sequence                : 0
  LBMs sent                                    : 0
```

 Copyright © 2013, Juniper Networks, Inc.

```
Valid in-order LBRs received                  : 0
Valid out-of-order LBRs received              : 0
LBRs received with corrupted data             : 0
LBRs sent                                     : 0
LTMs sent                                     : 0
LTMs received                                 : 0
LTRs sent                                     : 0
LTRs received                                 : 0
Sequence number of next LTM request           : 0
1DMs sent                                     : 0
Valid 1DMs received                           : 10
Invalid 1DMs received                         : 0
DMMs sent                                     : 0
DMRs sent                                     : 0
Valid DMRs received                           : 0
Invalid DMRs received                         : 0

Remote MEP identifier: 101
Remote MAC address: 00:90:69:0a:48:57
  Delay measurement statistics:
  Index  One-way delay  Two-way delay
           (usec)          (usec)
    1        370
    2        357
    3        344
    4        332
    5        319
    6        306
    7        294
    8        281
    9        269
    10       255
  Average one-way delay          : 312 usec
  Average one-way delay variation: 11 usec
  Best case one-way delay        : 255 usec
  Worst case one-way delay       : 370 usec
```

Related Documentation

- Guidelines for Configuring Routers to Support an ETH-DM Session on page 326
- Guidelines for Starting an ETH-DM Session on page 328
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 330
- On-Demand Mode on page 311
- *Ethernet Interfaces*

Configuring the Failure Notification Protocol

This topic describes how to configure the Ethernet Operations, Administration, and Maintenance (OAM) Failure Notification Protocol (FNP) on MX Series routers. The FNP detects link failures in a Carrier Ethernet network and broadcasts FNP messages when a failure occurs to all nodes affected by the link failure. To configure FNP functionality, include the **fnp** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam]
ethernet {
  fnp {
```

```
        interval <100ms | 1s | 10s | 1m | 10m>;
        loss-threshold number
        interface interface name {
            domain-id domain-id
        }
    }
}
```

The **interval** statement specifies the time between the transmission of FNP messages. You can specify 10 minutes (10m), 1 minute (1m), 10 seconds (10s), 1 second (1s), and 100 milliseconds (100ms). The **loss-threshold** statement specifies how many FNP messages can be lost before the FNP message is considered aged out and flushed. You must include the **interface *interface-name*** statement with the **domain-id *domain-id*** statement. The **domain-id** statement specifies a domain ID for the route. FNP messages can be received and processed on MX Series routers, but generating FNP messages is not supported.

The **show oam ethernet fnp interface**, **show oam ethernet fnp status**, and **show oam ethernet fnp messages** operational commands display the configured information.

FNP can be enabled only on logical interfaces that are part of a VPLS routing instance, and none of the logical interfaces in the VPLS routing instance should have CCM configured. FNP can be enabled on only one logical interface per physical interface.

Related Documentation

- *connectivity-fault-management*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 15

Configuring IEEE 802.1x Port-Based Network Access Control

IEEE 802.1x Port-Based Network Access Control Overview

MX Series routers support the IEEE 802.1x Port-Based Network Access Control (dot1x) protocol on Ethernet interfaces for validation of client and user credentials to prevent unauthorized access to a specified router port. Before authentication is complete, only 802.1x control packets are allowed and forwarded to the router control plane for processing. All other packets are dropped.

Authentication methods used must be 802.1x compliant. Authentication using RADIUS and Microsoft Active Directory servers is supported. The following user/client authentication methods are allowed:

- EAP-MD5 (RFC 3748)
- EAP-TTLS requires a server certificate (RFC 2716)
- EAP-TLS requires a client and server certificate
- PEAP requires only a server certificate

You can use both client and server certificates in all types of authentication except EAP-MD5.

NOTE: On the MX Series router, 802.1x can be enabled on bridged ports only and not on routed ports.

Dynamic changes to a user session are supported to allow the router administrator to terminate an already authenticated session by using the “RADIUS disconnect” message defined in RFC 3576.

Related Documentation

- Understanding the Administrative State of the Authenticator Port on page 368
- Understanding the Administrative Mode of the Authenticator Port on page 368
- Configuring the Authenticator on page 369
- Viewing the dot1x Configuration on page 369
- *Ethernet Interfaces*

Understanding the Administrative State of the Authenticator Port

The administrative state of an authenticator port can take any of the following three states:

- Force authorized—Allows network access to all users of the port without requiring them to be authenticated. This is equivalent to not having any authentication enabled on the port.
- Force unauthorized—Denies network access to all users of the port. This is equivalent to disabling the port.
- Automatic—This is the default mode where the authentication server response determines if the port is opened for traffic or not. Only the successfully authenticated clients are allowed access, all others are denied.

In Junos OS, the default mode is “automatic.” The “force authorized” and “force unauthorized” admin modes are not supported. You can achieve the functionality of “force authorized” mode by disabling **dot1x** on the required port. You can achieve the functionality of “force unauthorized” mode by disabling the port itself.

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 367
- Understanding the Administrative Mode of the Authenticator Port on page 368
- Configuring the Authenticator on page 369
- Viewing the dot1x Configuration on page 369
- *Ethernet Interfaces*

Understanding the Administrative Mode of the Authenticator Port

Junos OS supports the supplicant mode “single” and not the “single secure” nor “multiple” modes. The “Single” mode option authenticates only the first client that connects to a port. All other clients that connect later (802.1x compliant or noncompliant) are allowed free access on that port without any further authentication. If the first authenticated client logs out, all other users are locked out until a client authenticates again.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 367
- Understanding the Administrative State of the Authenticator Port on page 368
- Configuring the Authenticator on page 369
- Viewing the dot1x Configuration on page 369
- *Ethernet Interfaces*

Configuring the Authenticator

To configure the IEEE 802.1x Port-Based Network Access Control protocol on Ethernet interfaces you must configure the **authenticator** statement at the **[edit protocols dot1x]** hierarchy level. Use the **authentication-profile-name *access-profile-name*** statement to specify the authenticating RADIUS server, and use the **interface** statement to specify and configure the Gigabit Ethernet or Fast Ethernet interface on the router specifically for IEEE 802.1x protocol use; both at the **[edit protocols dot1x authenticator]** hierarchy level.

```
[edit protocols dot1x]
authenticator {
  authentication-profile-name access-profile-name;
  interface (xe-fpc/pic/port | ge-fpc/pic/port | fe-fpc/pic/port) {
    maximum-requests seconds;
    quiet-period seconds;
    reauthentication (disable | interval seconds);
    retries integer;
    server-timeout seconds;
    supplicant (single);
    supplicant-timeout seconds;
    transmit-period seconds;
  }
}
```

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 367
- Understanding the Administrative State of the Authenticator Port on page 368
- Understanding the Administrative Mode of the Authenticator Port on page 368
- Viewing the dot1x Configuration on page 369
- *Ethernet Interfaces*

Viewing the dot1x Configuration

Purpose To review and verify the dot1x configuration.

Action To view all **dot1x** configurations, use the **show dot1x interface** operational mode command. To view a **dot1x** configuration for a specific interface, use the **show dot1x interface (xe-*fpc/pic/port* | ge-*fpc/pic/port* | fe-*fpc/pic/port*) detail** operational mode command. See the *Network Interfaces Command Reference* for more information about this command.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- IEEE 802.1x Port-Based Network Access Control Overview on page 367
- Understanding the Administrative State of the Authenticator Port on page 368
- Understanding the Administrative Mode of the Authenticator Port on page 368
- Configuring the Authenticator on page 369
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 16

Configuring IEEE 802.3ah OAM Link-Fault Management

IEEE 802.3ah OAM Link-Fault Management Overview

Ethernet interfaces capable of running at 100 Mbps or faster on EX Series switches, MX Series, M Series (except M5 and M10 routers), and T Series routers support the IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology. Junos OS supports IEEE 802.3ah link-fault management.

The features of link-fault management are:

- Discovery
- Link monitoring
- Remote fault detection
- Remote loopback

The following features are not supported:

- Ethernet running on top of a Layer 2 protocol, such as Ethernet over ATM, is not supported in OAM configurations.
- Remote loopback is not supported on the 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The remote loopback feature mentioned in section 57.2.11 of IEEE 802.3ah is not supported on T4000 routers.

NOTE: **Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.**

Related Documentation

- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Configuring IEEE 802.3ah OAM Link-Fault Management

You can configure threshold values for fault events that trigger the sending of link event TLVs when the values exceed the threshold. To set threshold values for fault events on an interface, include the **event-thresholds** statement at the **[edit protocols oam ethernet link-fault-management interface]** hierarchy level.

You can also configure OAM threshold values within an action profile and apply the action profile to multiple interfaces. To create an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level.

You can configure Ethernet OAM either on an aggregate interface or on each of its member links. However, we recommend that you configure Ethernet OAM on the aggregate interface, and this will internally enable Ethernet OAM on the member links.

To view OAM statistics, use the **show oam ethernet link-fault-management** operational mode command. To clear OAM statistics, use the **clear oam ethernet link-fault-management statistics** operational mode command. To clear link-fault management state information and restart the link discovery process on Ethernet interfaces, use the **clear oam ethernet link-fault-management state** operational mode command. For more information about these commands, see the *Junos OS Operational Mode Commands*.

Related Documentation

- *event-thresholds*
- *action-profile*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387

- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Enabling IEEE 802.3ah OAM Support

To enable IEEE 802.3ah OAM support, include the **interface** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

When you enable IEEE 802.3ah OAM on a physical interface, the discovery process is automatically triggered.

Related Documentation

- **link-fault-management on page 591**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Configuring Link Discovery

When the IEEE 802.3ah OAM protocol is enabled on a physical interface, the discovery process is automatically triggered. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard.

Copyright © 2013, Juniper Networks, Inc.

You can specify the discovery mode used for IEEE 802.3ah OAM support. The discovery process is triggered automatically when OAM IEEE 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

To configure the discovery mode, include the **link-discovery** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
link-discovery (active | passive);
```

In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery.

Related Documentation

- **link-discovery on page 590**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Configuring the OAM PDU Interval

Periodic OAM PDUs are sent to perform link monitoring.

You can specify the periodic OAM PDU sending interval for fault detection.

To configure the sending interval, include the **pdu-interval** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
[edit protocol oam ethernet link-fault-management interface interface-name]
pdu-interval interval;
```

The periodic OAM PDU interval range is from 100 through 1000 milliseconds. The default sending interval is 1000 milliseconds.

Related Documentation

- **pdu-interval on page 615**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Configuring the OAM PDU Threshold

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

To configure the number of PDUs that can be missed from the peer, include the **pdu-threshold** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
pdu-threshold threshold-value;
```

The threshold value range is from 3 through 10. The default is three PDUs.

Related Documentation

- **pdu-threshold on page 616**

 Copyright © 2013, Juniper Networks, Inc.

- *Ethernet Interfaces*

Configuring Threshold Values for Local Fault Events on an Interface

You can configure threshold values on an interface for the local errors that trigger the sending of link event TLVs.

To set the error threshold values for sending event TLVs, include the **frame-error**, **frame-period**, **frame-period-summary**, and **symbol-period** statements at the **[edit protocols oam ethernet link-fault-management interface *interface-name* event-thresholds]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
event-thresholds {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Related Documentation

- *event-thresholds*
- *frame-error*
- *frame-period*

- *frame-period-summary*
- *symbol-period*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Disabling the Sending of Link Event TLVs

You can disable the sending of link event TLVs.

To disable the monitoring and sending of PDUs containing link event TLVs in periodic PDUs, include the **no-allow-link-events** statement at the **[edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
  negotiation-options]
no-allow-link-events;
```

Related Documentation

- *no-allow-link-events*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374

 Copyright © 2013, Juniper Networks, Inc.

- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Detecting Remote Faults

Fault detection is either based on flags or fault event type, length, and values (TLVs) received in OAM protocol data units (PDUs). Flags that trigger a link fault are:

- Critical Event
- Dying Gasp
- Link Fault

The link event TLVs are sent by the remote DTE by means of event notification PDUs. Link event TLVs are:

- Errored Symbol Period Event
- Errored Frame Event
- Errored Frame Period Event
- Errored Frame Seconds Summary Event

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375

Copyright © 2013, Juniper Networks, Inc.

Configuring an OAM Action Profile

You can create an action profile to define event fault flags and thresholds and the action to be taken. You can then apply the action profile to one or more interfaces.

To configure an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

```
action-profile profile-name {
  action {
    syslog;
    link-down;
    send-critical-event;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    protocol-down;
  }
}
```

Related Documentation

 Copyright © 2013, Juniper Networks, Inc.

Specifying the Actions to Be Taken for Link-Fault Management Events

You can specify the action to be taken by the system when the configured link-fault event occurs. Multiple action profiles can be applied to a single interface. For each action-profile, at least one event and one action must be specified. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all the actions are executed.

You might want to set a lower threshold for a specific action such as logging the error and set a higher threshold for another action such as sending a critical event TLV.

To specify the action, include the **action** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
event {
  link-adjacency-loss;
  protocol-down;
}
action {
  syslog;
  link-down;
  send-critical-event;
}
```

To create a system log entry when the link-fault event occurs, include the **syslog** statement.

To administratively disable the link when the link-fault event occurs, include the **link-down** statement.

To send IEEE 802.3ah link event TLVs in the OAM PDU when a link-fault event occurs, include the **send-critical-event** statement.

NOTE: If multiple actions are specified in the action profile, all of the actions are executed in no particular order.

Related Documentation

- *action*
- *syslog*
- *link-down*
- *send-critical-event*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Monitoring the Loss of Link Adjacency

You can specify actions be taken when link adjacency is lost. When link adjacency is lost, the system takes the action defined in the **action** statement of the action profile.

To configure the system to take action when link adjacency is lost, include the **link-adjacency-loss** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
link-adjacency-loss;
```

Related Documentation

- *link-adjacency-loss*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Monitoring Protocol Status

The CCC-DOWN flag is associated with a circuit cross-connect (CCC) connection, Layer 2 circuit, and Layer 2 VPN, which send the CCC-DOWN status to the kernel. The CCC-DOWN flag indicates that the CCC is down. The CCC-DOWN status is sent to the kernel when the CCC connection, Layer 2 circuit, or Layer 2 VPN is down. This in turn, brings down the CE-facing PE interface associated with the CCC connection, Layer 2 circuit, or Layer 2 VPN.

Copyright © 2013, Juniper Networks, Inc.

When the CCC-DOWN flag is signaled to the IEEE 802.3ah protocol, the system takes the action defined in the **action** statement of the action profile. For additional information about Layer 2 circuits, see the Junos OS Layer 2 Circuits Feature Guide, Junos OS VPNs Configuration Guide.

To monitor the IEEE 802.3ah protocol, on the CE-facing PE interface, include the **protocol-down** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
protocol-down;
```

NOTE: If multiple events are specified in the action profile, all the events must occur before the specified action is taken.

Related Documentation

- **protocol-down on page 633**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Configuring Threshold Values for Fault Events in an Action Profile

You can configure link event thresholds for received error events that trigger the action specified in the **action** statement. You can then apply the action profile to one or more interfaces.

To configure link event thresholds, include the **link-event-rate** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Related Documentation

- *link-event-rate*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Applying an Action Profile

You can apply an action profile to one or more interfaces.

Copyright © 2013, Juniper Networks, Inc.

To apply an action profile to an interface, include the **apply-action-profile** statement at the **[edit protocols oam ethernet link-fault-management action-profile interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
apply-action-profile profile-name;
```

Related Documentation

- *apply-action-profile*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Setting a Remote Interface into Loopback Mode

You can configure the software to set the remote DTE into loopback mode on the following interfaces:

- IQ2 and IQ2-E Gigabit Ethernet interfaces
- Ethernet interfaces on the MX Series routers or EX Series switches

Junos OS can place a remote DTE into loopback mode (if remote-loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote-loopback request and puts the interface into remote-loopback mode. When the interface is in remote-loopback mode, all frames

 Copyright © 2013, Juniper Networks, Inc.

except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent to the management plane and processed.

To configure remote loopback, include the **remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
remote-loopback;
```

To take the remote DTE out of loopback mode, remove the **remote-loopback** statement from the configuration.

Related Documentation

- *remote-loopback*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Enabling Remote Loopback Support on the Local Interface on page 387
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Enabling Remote Loopback Support on the Local Interface

You can allow a remote DTE to set a local interface into remote loopback mode on IQ2 and IQ2-E Gigabit Ethernet interfaces and all Ethernet interfaces on the MX Series routers and EX Series switches. When a remote-loopback request is sent by a remote DTE, the Junos OS places the local interface into loopback mode. When an interface is in loopback mode, all frames except OAM PDUs are looped back without any changes to the frames. OAM PDUs continue to be sent to the management plane and processed. By default, the remote loopback feature is not enabled.

Copyright © 2013, Juniper Networks, Inc.

To enable remote loopback, include the **allow-remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
  negotiation-options]
allow-remote-loopback;
```

NOTE: Activation of OAM remote loopback may result in data frame loss.

Related Documentation

- *allow-remote-loopback*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 388
- *Ethernet Interfaces*

Example: Configuring IEEE 802.3ah OAM Support on an Interface

Configure 802.3ah OAM support on a 10-Gigabit Ethernet interface:

```
[edit]
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface xe-0/0/0 {
```

 Copyright © 2013, Juniper Networks, Inc.

```
                    link-discovery active;
                    pdu-interval 800;
                    pdu-threshold 4;
                    remote-loopback;
                    negotiation-options {
                      allow-remote-loopback;
                    }
                    event-thresholds {
                      frame-error 30;
                      frame-period 50;
                      frame-period summary 40;
                      symbol-period 20;
                    }
                  }
                }
              }
            }
          }
```

Related Documentation

- **link-fault-management on page 591**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Enabling IEEE 802.3ah OAM Support on page 374
- Configuring Link Discovery on page 374
- Configuring the OAM PDU Interval on page 375
- Configuring the OAM PDU Threshold on page 376
- Configuring Threshold Values for Local Fault Events on an Interface on page 377
- Disabling the Sending of Link Event TLVs on page 378
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 381
- Monitoring the Loss of Link Adjacency on page 383
- Monitoring Protocol Status on page 383
- Configuring Threshold Values for Fault Events in an Action Profile on page 385
- Applying an Action Profile on page 385
- Setting a Remote Interface into Loopback Mode on page 386
- Enabling Remote Loopback Support on the Local Interface on page 387
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 17

Configuring VRRP and VRRP for IPv6

VRRP and VRRP for IPv6 Overview

You can configure the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6 for the following interfaces:

- Ethernet
- Fast Ethernet
- Tri-Rate Ethernet copper
- Gigabit Ethernet
- 10-Gigabit Ethernet LAN/WAN PIC
- Ethernet logical interfaces

VRRP and VRRP for IPv6 allow hosts on a LAN to make use of redundant routers on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routers share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routers is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master router, thus always providing a virtual default router and allowing traffic on the LAN to be routed without relying on a single router.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

For VRRP and VRRP for IPv6 overview information, configuration guidelines, and statement summaries, see the *Junos OS High Availability Library for Routing Devices*.

Related Documentation

- Configuring VRRP and VRRP for IPv6 on page 392

- *Ethernet Interfaces*

Configuring VRRP and VRRP for IPv6

To configure VRRP or VRRP for IPv6, include the **vrrp-group** or **vrrp-inet6-group** statement, respectively. These statements are available at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

The VRRP and VRRP IPv6 configuration statements are as follows:

```
(vrrp-group | vrrp-inet-group) group-number {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority-number number;
  track {
    priority-hold-time;
    interface interface-name {
      priority-cost priority;
      bandwidth-threshold bits-per-second {
        priority-cost;
      }
    }
  }
  virtual-address [ addresses ];
}
```

You can configure VRRP IPv6 with a global unicast address.

To trace VRRP and VRRP for IPv6 operations, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  file <filename> <files number <match regular-expression <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

When there are multiple VRRP groups, there is a few seconds delay between the time the first gratuitous ARP is sent out and the rest of the gratuitous ARP are sent. Configuring failover-delay compensates for this delay. To configure the failover delay from 500 to 2000 milliseconds for VRRP and VRRP for IPv6 operations, include the **failover-delay *milliseconds*** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
failover-delay milliseconds;
```


Copyright © 2013, Juniper Networks, Inc.

To configure the startup period for VRRP and VRRP for IPv6 operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
startup-silent-period seconds;
```

To enable VRRPv3, set the **version-3** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
version-3;
```

Related Documentation

- *VRRP and VRRP for IPv6 Feature Guide for Routing Devices*
- *failover-delay*
- *traceoptions*
- *failover-delay*
- *vrrp-group*
- VRRP and VRRP for IPv6 Overview on page 391
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 18

Configuring Gigabit Ethernet Accounting and Policing

Gigabit Ethernet Accounting and Policing Overview

For Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure granular per-VLAN class-of-service (CoS) capabilities and extensive instrumentation and diagnostics on a per-VLAN and per-MAC address basis.

VLAN rewrite, tagging, and deleting enables you to use VLAN address space to support more customers and services.

VPLS allows you to provide a point-to-multipoint LAN between a set of sites in a VPN. Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) are combined with VPLS to deliver metro Ethernet service.

For Gigabit Ethernet IQ2 and IQ2-E and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, you can apply Layer 2 policing to logical interfaces in the egress or ingress direction. Layer 2 policers are configured at the **[edit firewall]** hierarchy level. You can also control the rate of traffic sent or received on an interface by configuring a policer overhead at the **[edit chassis fpc *slot-number* pic *slot-number*]** hierarchy level.

Table 29 on page 396 lists the capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Copyright © 2013, Juniper Networks, Inc.

Table 29: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
Layer 2		
802.3ad link aggregation	Yes	Yes
Maximum VLANs per port	384	1023
Maximum transmission unit (MTU) size	9192	9192
MAC learning	Yes	Yes
MAC accounting	Yes	Yes
MAC filtering	Yes	Yes
Destinations per port	960	960
Sources per port	64	64
Hierarchical MAC policers	Yes, premium and aggregate	No, aggregate only
Multiple TPID support and IP service for nonstandard TPIDs	Yes	Yes
Multiple Ethernet encapsulations	Yes	Yes
Dual VLAN tags	Yes	No
VLAN rewrite	Yes	No
Layer 2 VPNs		
VLAN CCC	Yes	Yes
Port-based CCC	Yes	Yes
Extended VLAN CCC Virtual Metropolitan Area Network (VMAN) Tag Protocol	Yes	Yes
CoS		
PIC-based egress queues	Yes	Yes
Queued VLANs	Yes	No
VPLS	Yes	Yes

Copyright © 2013, Juniper Networks, Inc.

For more information about configuring VPLS, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Feature Guides*.

You can also configure CoS on logical IQ interfaces. For more information, see the *Junos OS Class of Service Library for Routing Devices*.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 403
- Configuring MAC Address Accounting on page 406

- *Configuring a Policer Overhead*
- *Ethernet Interfaces*

Configuring Gigabit Ethernet Policers

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can define rate limits for premium and aggregate traffic received on the interface. These policers allow you to perform simple traffic policing without configuring a firewall filter. First you configure the Ethernet policer profile, next you classify ingress and egress traffic, then you can apply the policer to a logical interface.

For Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the policer rates you configure can be different than the rates on the Packet Forward Engine. The difference results from Layer 2 overhead. The PIC accounts for this difference.

NOTE:
On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

For information on understanding Ethernet frame statistics, see the *MX Series Layer 2 Configuration Guide*.

This section contains the following topics:

- Configuring a Policer on page 398
- Specifying an Input Priority Map on page 398
- Specifying an Output Priority Map on page 399

- Applying a Policer on page 400
- Configuring MAC Address Filtering on page 401
- Example: Configuring Gigabit Ethernet Policers on page 402

Configuring a Policer

To configure an Ethernet policer profile, include the **ethernet-policer-profile** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
ethernet-policer-profile {
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
}
```

In the Ethernet policer profile, the aggregate-priority policer is mandatory; the premium-priority policer is optional.

For aggregate and premium policers, you specify the bandwidth limit in bits per second. You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 4.29 Gbps.

The maximum burst size controls the amount of traffic bursting allowed. To determine the burst-size limit, you can multiply the bandwidth of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur:

burst size = *bandwidth* x *allowable time for burst traffic*

If you do not know the interface bandwidth, you can multiply the maximum MTU of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. The burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 MB.

Specifying an Input Priority Map

An input priority map identifies ingress traffic with specified IEEE 802.1p priority values, and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an input priority map by including the **ieee802.1 premium** statement at the **[edit interfaces *interface-name* gigether-options ethernet-policer-profile input-priority-map]** hierarchy level:

 Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]
ieee802.1p premium [ values ];
```

The priority values can be from 0 through 7. The remaining traffic is classified as nonpremium (or aggregate). For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 402.

NOTE: On IQ2 and IQ2-E interfaces and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

Specifying an Output Priority Map

An output priority map identifies egress traffic with specified queue classification and packet loss priority (PLP), and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an output priority map by including the **classifier** statement at the **[edit interfaces *interface-name* gigether-options ethernet-policer-profile output-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile
  output-priority-map]
classifier {
  premium {
    forwarding-class class-name {
      loss-priority (high | low);
    }
  }
}
```

You can define a forwarding class, or you can use a predefined forwarding class. Table 30 on page 399 shows the predefined forwarding classes and their associated queue assignments.

Table 30: Default Forwarding Classes

Forwarding Class Name	Queue
best-effort	Queue 0
expedited-forwarding	Queue 1
assured-forwarding	Queue 2
network-control	Queue 3

For more information about CoS forwarding classes, see the *Junos OS Class of Service Library for Routing Devices*. For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 402.

Applying a Policer

On all MX Series Router interfaces, Gigabit Ethernet IQ, IQ2, and IQ2-E PICs, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can apply input and output policers that define rate limits for premium and aggregate traffic received on the logical interface. Aggregate policers are supported on Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

These policers allow you to perform simple traffic policing without configuring a firewall filter. For information about defining these policers, see “Configuring Gigabit Ethernet Policers” on page 397.

To apply policers to specific source MAC addresses, include the **accept-source-mac** statement:

```
accept-source-mac {
  mac-address mac-address {
    policer {
      input cos-policer-name;
      output cos-policer-name;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where ***n*** is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include multiple **mac-address** statements in the logical interface configuration.

NOTE: On untagged Gigabit Ethernet interfaces you should not configure the **source-address-filter** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces you should not configure the **source-address-filter** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options unit *logical-unit-number*]** hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

 Copyright © 2013, Juniper Networks, Inc.

NOTE: If the remote Ethernet card is changed, the interface does not accept traffic from the new card because the new card has a different MAC address.

The MAC addresses you include in the configuration are entered into the router's MAC database. To view the router's MAC database, enter the **show interfaces mac-database *interface-name*** command:

user@host> **show interfaces mac-database *interface-name***

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.

NOTE: On IQ2 and IQ2-E PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

You can use the same policer one or more times.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

Configuring MAC Address Filtering

You cannot explicitly define traffic with specific source MAC addresses to be rejected; however, for Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet DPCs on MX Series routers, you can block all incoming packets that do not have a source address specified in the **accept-source-mac** statement. For more information about the **accept-source-mac** statement, see "Applying a Policer" on page 400.

To enable this blocking, include the **source-filtering** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level:

```
[edit interfaces interface-name gigether-options]
source-filtering;
```

For more information about the **source-filtering** statement, see "Enabling Ethernet MAC Address Filtering" on page 42.

To accept traffic even though it does not have a source address specified in the **accept-source-mac** statement, include the **no-source-filtering** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level:

```
[edit interfaces interface-name gigether-options]
no-source-filtering;
```

Copyright © 2013, Juniper Networks, Inc.

For more information about the **accept-source-mac** statement, see "Applying a Policer" on page 400.

Example: Configuring Gigabit Ethernet Policers

Configure interface **ge-6/0/0** to treat priority values 2 and 3 as premium. On ingress, this means that IEEE 802.1p priority values **2** and **3** are treated as premium. On egress, it means traffic that is classified into queue 0 or 1 with PLP of low and queue 2 or 3 with PLP of high, is treated as premium.

Define a policer that limits the premium bandwidth to 100 Mbps and burst size to 3 k, and the aggregate bandwidth to 200 Mbps and burst size to 3 k.

Specify that frames received from the MAC address **00:01:02:03:04:05** and the VLAN ID **600** are subject to the policer on input and output. On input, this means frames received with the source MAC address **00:01:02:03:04:05** and the VLAN ID 600 are subject to the policer. On output, this means frames transmitted from the router with the destination MAC address **00:01:02:03:04:05** and the VLAN ID **600** are subject to the policer.

```
[edit interfaces]
ge-6/0/0 {
  gigether-options {
    ether-switch-profile {
      ether-policer-profile {
        input-priority-map {
          ieee-802.1p {
            premium [ 2 3 ];
          }
        }
        output-priority-map {
          classifier {
            premium {
              forwarding-class best-effort {
                loss-priority low;
              }
              forwarding-class expedited-forwarding {
                loss-priority low;
              }
              forwarding-class assured-forwarding {
                loss-priority high;
              }
              forwarding-class network-control {
                loss-priority high;
              }
            }
          }
        }
        policer policer-1 {
          premium {
            bandwidth-limit 100m;
            burst-size-limit 3k;
          }
          aggregate {
```

 Copyright © 2013, Juniper Networks, Inc.

```
                    bandwidth-limit 200m;
                    burst-size-limit 3k;
                  }
                }
              }
            }
          }
          unit 0 {
            accept-source-mac {
              mac-address 00:01:02:03:04:05 {
                policer {
                  input policer-1;
                  output policer-1;
                }
              }
            }
          }
        }
```

Related Documentation

- Gigabit Ethernet Accounting and Policing Overview on page 395
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 403
- Configuring MAC Address Accounting on page 406

- *Configuring a Policer Overhead*
- *Ethernet Interfaces*

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

For Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers, you can configure two-color and tricolor marking policers and apply them to logical interfaces to prevent traffic on the interface from consuming bandwidth inappropriately.

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks router architectures support three types of policer:

- Two-color policer—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit in some way, or simply discard them. A policer is most useful for metering traffic at the port (physical interface) level.

Copyright © 2013, Juniper Networks, Inc.

- Single-rate tricolor marking (srTCM)—A single-rate tricolor marking policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets arriving are below the CBS (green), exceed the CBS (yellow) but not the EBS, or exceed the EBS (red). Single-rate TCM is most useful when a service is structured according to packet length and not peak arrival rate.
- Two-rate Tricolor Marking (trTCM)—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and peak information rate (PIR), along with their associated burst sizes, the CBS and EBS. Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets arriving are below the CIR (green), exceed the CIR (yellow) but not the PIR, or exceed the PIR (red). Two-rate TCM is most useful when a service is structured according to arrival rates and not necessarily packet length.

Unlike policing (described in "Configuring Gigabit Ethernet Policers" on page 397), configuring two-color policers and tricolor marking policers requires that you configure a firewall filter.

This section contains the following topics:

Configuring a Policer

Two-color and tricolor marking policers are configured at the **[edit firewall]** hierarchy level.

A tricolor marking policer polices traffic on the basis of metering rates, including the CIR, the PIR, their associated burst sizes, and any policing actions configured for the traffic.

To configure tricolor policer marking, include the **three-color-policer** statement with options at the **[edit firewall]** hierarchy level:

```
[edit firewall]
three-color-policer name {
  action {
    loss-priority high {
      then discard;
    }
  }
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
```


Copyright © 2013, Juniper Networks, Inc.

```
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
    }
}
```

For more information about configuring tricolor policer markings, see the *Routing Policy Feature Guide for Routing Devices* and the *Junos OS Class of Service Library for Routing Devices*.

Applying a Policer

Apply a two-color policer or tricolor policer to a logical interface to prevent traffic on the interface from consuming bandwidth inappropriately. To apply two-color or tricolor policers, include the **layer2-policer** statement:

```
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    policer-name;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Use the **input-policer** statement to apply a two-color policer to received packets on a logical interface and the **input-three-color** statement to apply a tricolor policer. Use the **output-policer** statement to apply a two-color policer to transmitted packets on a logical interface and the **output-three-color** statement to apply a tricolor policer. The specified policers must be configured at the **[edit firewall]** hierarchy level. For each interface, you can configure a three-color policer or two-color input policer or output policers—you cannot configure both a three-color policer and a two-color policer.

Example: Configuring and Applying a Policer

Configure tricolor policers and apply them to an interface:

```
[edit firewall]
three-color-policer three-color-policer-color-blind {
    logical-interface-policer;
    two-rate {
        color-blind;
        committed-information-rate 1500000;
        committed-burst-size 150;
        peak-information-rate 3;
        peak-burst-size 300;
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  three-color-policer three-color-policer-color-aware {
    logical-interface-policer;
    two-rate {
      color-aware;
      committed-information-rate 1500000;
      committed-burst-size 150;
      peak-information-rate 3;
      peak-burst-size 300;
    }
  }
  [edit interfaces ge-1/1/0]
  unit 1 {
    layer2-policer {
      input-three-color three-color-policer-color-blind;
      output-three-color three-color-policer-color-aware;
    }
  }
```

Configure a two-color policer and apply it to an interface:

```
[edit firewall]
policer two-color-policer {
  logical-interface-policer;
  if-exceeding {
    bandwidth-percent 90;
    burst-size-limit 300;
  }
  then loss-priority-high;
}
[edit interfaces ge-1/1/0]
unit 2 {
  layer2-policer {
    input-policer two-color-policer;
    output-policer two-color-policer;
  }
}
```

Related Documentation

- Gigabit Ethernet Accounting and Policing Overview on page 395
- Configuring Gigabit Ethernet Policers on page 397
- Configuring MAC Address Accounting on page 406
- *Configuring a Policer Overhead*
- *Ethernet Interfaces*

Configuring MAC Address Accounting

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet DPCs on MX Series routers, you can configure whether source and destination MAC addresses are dynamically learned. To configure MAC address accounting, include

 Copyright © 2013, Juniper Networks, Inc.

the **mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
mac-learn-enable;
```

To prohibit the interface from dynamically learning source and destination MAC addresses, include the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
no-mac-learn-enable;
```

MAC address learning is based on source addresses. You can start accounting for traffic after it has been sent from the MAC address. Once the MAC address is learned, the frames and bytes transmitted to or received from the MAC address can be tracked.

NOTE: DPCs and MPCs support MAC address accounting. DPCs support both source and destination MAC address accounting. MPCs support only source MAC address accounting.

Related Documentation

- Gigabit Ethernet Accounting and Policing Overview on page 395
- Configuring Gigabit Ethernet Policers on page 397
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 403
- *Configuring a Policer Overhead*
- *Ethernet Interfaces*

Accounting of the Layer 2 Overhead Attribute in Interface Statistics

On MX Series routers, you can configure the physical interface and logical interface statistics to include the Layer 2 overhead size (header and trailer bytes) for both ingress and egress interfaces. Both the transit and total statistical information are computed and displayed for each logical interface. This functionality is supported on 1-Gigabit and 10-Gigabit Ethernet interfaces on Dense Port Concentrators (DPCs) and Modular Port Concentrators (MPCs). Also, this feature is supported on 10-Gigabit Ethernet interfaces on MX Series routers with MPC4E. To enable the Layer 2 overhead bytes to be counted in the interface statistics at the PIC level, you must use the **account-layer2-overhead** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level.

You can also enable the Layer 2 overhead bytes for computation in the logical interface statistics by configuring the **account-layer2-overhead (*value* | <ingress *bytes* | egress *bytes*>)** statement at the **[edit interface *interface-name* unit *logical-unit-number*]** hierarchy level. If you configure this capability, all the Layer 2 header details (Layer 2 header and cyclic redundancy check [CRC]) based on the Layer 2 encapsulation configured for an interface are calculated and displayed in the physical and logical interface statistics for ingress and egress interfaces in the output of the **show interfaces *interface-name*** commands. For physical and logical interfaces, the **Input bytes** and **Output bytes** fields

Copyright © 2013, Juniper Networks, Inc.

under the Traffic statistics section in the output of the **show interfaces** ***interface-name*** **<detail | extensive>** command include the Layer 2 overhead of the packets. For physical and logical interfaces, the Input rate and Output rate fields under the Traffic statistics section in the output of the **show interfaces** ***interface-name*** **<media | statistics>** command include the Layer 2 overhead of the packets. For logical interfaces, the values for the newly added **Egress accounting overhead** and **Ingress accounting overhead** fields display the Layer 2 overhead size for transmitted and received packets respectively.

The input and output octets at the physical or logical interface configured on the PIC includes all the Layer 2 headers. All the logical interfaces on the PIC, including the ae and the non-ae interfaces, are processed for Layer 2 overhead accounting for the arriving and exiting packets. This method of operation impacts the transit statistics that are primarily used for subscriber accounting and billing purposes in customer networks.

Table 31 on page 408 lists the adjustment bytes that are counted based on the encapsulation on the logical interface over the Ethernet interface, when you enable accounting of Layer 2 overhead in interface statistics at the PIC level. The values for the adjustment bytes that are listed for all types of encapsulation are the same for DPCs and MPCs, with the only exception being for the VLAN CCC adjustment value. On DPCs, the VLAN CCC adjustment value is –4 bytes and on MPCs, the VLAN CCC adjustment value is +4 bytes.

Table 31: Adjustment Bytes for Logical Interfaces over Ethernet Interfaces

Encapsulation Type on Logical Interfaces	Number of Adjustment Bytes	Description
Ethernet DIXv2 (IP datagrams over Ethernet)	18	Untagged (includes CRC)
Ethernet DIXv2 (IP datagrams over Ethernet)	22	Single-tagged (includes CRC)
Ethernet DIXv2 (IP datagrams over Ethernet)	26	Double-tagged (includes CRC)
VLAN Bridge	4	CRC
VLAN CCC	4	CRC
VLAN TCC	18	Untagged (includes CRC)
VLAN TCC	22	Single-tagged (includes CRC)
VLAN TCC	26	Double-tagged (includes CRC)
VLAN VPLS	4	CRC

Guidelines for Configuring the Computation of Layer 2 Overhead in Interface Statistics

Keep the following points in mind when you configure the feature to compute Layer 2 overhead in interface statistics:

 Copyright © 2013, Juniper Networks, Inc.

- When you configure a native VLAN ID on a logical interface, the Layer 2 header adjustment for input statistics is different for tagged and untagged packets. For such interfaces, if you configure the setting to account for Layer 2 overhead, incorrect statistics might be displayed.
- An untagged packet is considered as a tagged packet and an additional 4 bytes are appended to the counter values displayed in the output of the **show interface** command.
- The computed statistics might not be completely accurate in scenarios where the packets are dropped after they have been included in the interface statistics, but before the packets reach the destination.
- Label-switched interface (LSI) statistics on the ingress direction of interfaces do not include the Layer 2 overhead bytes because this functionality of accounting Layer 2 overhead is not supported for such LSI interfaces.
- Layer 2 overhead accounting is not supported for inline service (si) interfaces.
- The total statistics of physical interfaces do not indicate the complete Layer 2 adjusted statistics. This behavior occurs because the total statistics count is the sum of transit and local statistics. Only the transit statistics are adjusted for Layer 2 and the local statistics are not adjusted for Layer 2.
- Statistics on ae interfaces are calculated in the same manner as non-ae interfaces.
- Adjustment bytes are applicable only for transit statistics that are displayed for logical interfaces.
- For physical interfaces, the adjustment bytes for transit traffic and the non-adjusted bytes for local or protocol-specific traffic are combined and displayed in the output of the **show interfaces** command. (Segregation is not possible.)
- If you configure the **account-layer2-overhead (** ***value*** **| <ingress** ***bytes*** **| egress** ***bytes*>)** statement at the **[edit interface** ***interface-name*** **unit** ***logical-unit-number*]** hierarchy level, the Layer 2 overhead computation might not work correctly for logical interfaces over an aggregated Ethernet (ae) interface that is created across two different line cards having different forwarding hardware (such as MPCs and DPCs). This method of processing applies only for VLAN circuit cross-connect (CCC) encapsulation.
- The Layer 2 overhead bytes in interface statistics are saved across a unified ISSU or a graceful Routing Engine switchover (GRES) operation.

Related Documentation

- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 410
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 412
- **account-layer2-overhead (Interface Level) on page 550**
- **account-layer2-overhead (PIC Level) on page 551**
- *Ethernet Interfaces*

Configuring Layer 2 Overhead Accounting in Interface Statistics

This topic contains the following sections that describe the configuration of Layer 2 overhead accounting for interface statistics at the PIC level and logical interface level:

- Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level on page 410
- Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the Logical Interface Level on page 410

Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level

You can configure the **account-layer2-overhead** statement at the **edit chassis fpc *slot-number* pic *pic-number*** hierarchy level to enable accounting of Layer 2 overhead bytes in the ingress and egress interface statistics at the PIC level.

CAUTION: **If you modify the setting for accounting of Layer 2 overhead bytes at the PIC level, the PIC is rebooted, causing all of the physical and logical interfaces to be deleted and readded on the PIC. Due to this behavior, we recommend that you exercise caution while using this feature.**

The computation method of Layer 2 overhead on different interface types is as follows:

- For Ethernet interfaces, all the Layer 2 headers are counted.
- For non-Ethernet interfaces, the Frame Relay, PPP, or Cisco HDLC headers are counted, while the bit or byte stuffing headers are excluded.

To enable accounting of Layer 2 overhead at the PIC level for ingress and egress traffic on interfaces:

1. Access a DPC or an MPC-occupied slot and the PIC where the interface is to be enabled.

   [edit chassis]
   user@host# **edit fpc *slot-number* pic *number***

2. Specify the Layer 2 overhead value in bytes that is the octet adjustment per packet added to the total octet count for ingress and egress traffic on all the interfaces in the PIC.

   [edit chassis fpc *slot-number* pic *number*]
   user@host# **set account-layer2-overhead**

Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the Logical Interface Level

You can configure the **account-layer2-overhead (*value* | <ingress *bytes* | egress *bytes*>)** statement at the **[edit interface *interface-name* unit *logical-unit-number*]** hierarchy level to enable accounting of Layer 2 overhead bytes in the ingress and egress interface

 Copyright © 2013, Juniper Networks, Inc.

statistics. If you do not configure this setting at the logical interface level, the Layer 2 overhead value is assumed as 0.

Both the options of enabling Layer 2 overhead accounting at the PIC level and the logical interface level are disabled by default. You can configure either of the options or both the options at the same time. If you enable Layer 2 overhead accounting at both the PIC level and the logical interface level, the logical interface-level setting takes precedence over the PIC-level setting. When both the options are configured, if you configure only the **account-layer2-overhead ingress *bytes*** statement at the **[edit interface *interface-name* unit *logical-unit-number*]** hierarchy level, then the ingress Layer 2 overhead is counted with the explicit adjustment value you specify. Egress Layer 2 overhead is adjusted automatically based on the encapsulation on the logical interface. Similarly, if you configure the Layer 2 overhead accounting only at the egress side using this statement, the egress overhead is calculated using the value you specify, whereas the ingress overhead is adjusted automatically.

The counted Layer 2 overhead bytes are displayed in the interface statistics.

CAUTION: **If you modify the setting for accounting of Layer 2 overhead bytes at the interface level, the changed setting is effective on the packets that are received or transmitted from that time. The logical interface is deleted and readded. All the prior statistics are lost and the traffic is disrupted. Due to this behavior, we recommend that you exercise caution while using this feature.**

The overhead value option and the ingress or egress options are mutually exclusive. You can either specify the overhead value option to account for Layer 2 overhead in both the ingress and egress traffic, or the ingress and egress options separately to account for Layer 2 overhead to the ingress and egress traffic respectively. The Layer 2 overhead byte can be a number in the range –128 through +127.

To configure accounting of Layer 2 overhead at the logical interface level:

1. Access the logical interface.

   [edit interfaces]
   user@host# **edit *interface-name* unit *logical-unit-number***

2. (Optional; for both ingress and egress traffic) Specify the Layer 2 overhead value in bytes that is the octet adjustment per packet added to the total octet count for ingress and egress traffic on the logical interface.

   [edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# **set account-layer2-overhead 127**

   In this example, 127 bytes are added per packet for ingress and egress traffic on the logical interface.

Copyright © 2013, Juniper Networks, Inc.

NOTE: The **overhead value** ***bytes*** option and the **ingress** ***bytes*** or **egress** ***bytes*** options are mutually exclusive. You can either specify the **overhead value** ***bytes*** option to account for Layer 2 overhead in both the ingress and egress traffic, or the **ingress** ***bytes*** and **egress** ***bytes*** options separately to account for Layer 2 overhead to the ingress and egress traffic respectively.

3. (Optional; for ingress traffic only) Specify the Layer 2 overhead value in bytes that is the octet adjustment per packet added to the total octet count for ingress traffic on the logical interface.

   [edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# **set account-layer2-overhead overhead ingress 100**

   In this example, 100 bytes are added per packet for ingress traffic on the logical interface. For egress traffic, Layer 2 overhead is adjusted automatically based on the encapsulation on the logical interface.

4. (Optional; for egress traffic only) Specify the Layer 2 overhead value in bytes that is the octet adjustment per packet added to the total octet count for egress traffic on the logical interface.

   [edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# **set account-layer2-overhead overhead egress 100**

   In this example, 100 bytes are added per packet for egress traffic on the logical interface. For ingress traffic, Layer 2 overhead is adjusted automatically based on the encapsulation on the logical interface.

Verifying the Accounting of Layer 2 Overhead in Interface Statistics

Purpose Display information about the Layer 2 overhead bytes that are counted in interface statistics for egress and ingress traffic on Ethernet interfaces.

Action

- To display information about the Layer 2 overhead bytes that are counted in interface statistics:

NOTE: For physical and logical interfaces, the values displayed for the **Input rate** and **Output rate** fields under the Traffic statistics section include the Layer 2 overhead of the packets.

user@host> **show interfaces ge-5/2/0 statistics detail**

```
Physical interface: ge-5/2/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 519, Generation: 149
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
 Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
```

 Copyright © 2013, Juniper Networks, Inc.

```
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:61:d9:74, Hardware address: 00:1d:b5:61:d9:74
  Last flapped   : 2009-11-11 11:24:00 PST (09:23:08 ago)
  Statistics last cleared: 2009-11-11 17:50:58 PST (02:56:10 ago)
  Traffic statistics:
   Input  bytes  :             271524                    0 bps
   Output bytes  :           37769598                  352 bps
   Input  packets:               3664                    0 pps
   Output packets:             885790                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                 0
    Output bytes  :          16681118
    Input  packets:                 0
    Output packets:            362633
   Multicast statistics:
    IPV4 multicast statistics:
     Input  bytes  :             112048                    0 bps
     Output bytes  :           20779920                    0 bps
     Input  packets:               1801                    0 pps
     Output packets:             519498                    0 pps
    IPV6 multicast statistics:
     Input  bytes  :             156500                    0 bps
     Output bytes  :           16681118                    0 bps
     Input  packets:               1818                    0 pps
     Output packets:             362633                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets     Dropped packets

    0 best-effort             882558              882558
 0
    1 expedited-fo                 0                   0
 0
    2 assured-forw                 0                   0
 0
    3 network-cont              3232                3232
 0
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Traffic statistics:
     Input  bytes  :             271524
     Output bytes  :           37769598
     Input  packets:               3664
     Output packets:             885790
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :          16681118
      Input  packets:                 0
```

Copyright © 2013, Juniper Networks, Inc.

```
      Output packets:              362633
    Local statistics:
     Input  bytes  :               271524
     Output bytes  :               308560
     Input  packets:                 3664
     Output packets:                 3659
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :             37461038                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:               882131                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :            16681118
      Input  packets:                   0
      Output packets:              362633
    Multicast statistics:
     IPV4 multicast statistics:
      Input  bytes  :               112048                    0 bps
      Output bytes  :             20779920                    0 bps
      Input  packets:                 1801                    0 pps
      Output packets:               519498                    0 pps
     IPV6 multicast statistics:
      Input  bytes  :               156500                    0 bps
      Output bytes  :             16681118                    0 bps
      Input  packets:                 1818                    0 pps
      Output packets:               362633                    0 pps
    Protocol inet, MTU: 1500, Generation: 151, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 40.40.40.0/30, Local: 40.40.40.2, Broadcast: 40.40.40.3,
Generation: 167
    Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::40.40.40.0/126, Local: ::40.40.40.2
    Generation: 169
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974
    Protocol multiservice, MTU: Unlimited, Generation: 171
    Generation: 153, Route table: 0
      Policer: Input: __default_arp_policer__
```

For more information about the **show interfaces statistics** command, see *show interfaces statistics* in the *Junos OS Operational Mode Commands*.

Related Documentation

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics on page 407
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 410
- *show interfaces (Gigabit Ethernet)*
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 19

Configuring Gigabit Ethernet Autonegotiation

Gigabit Ethernet Autonegotiation Overview

Autonegotiation is enabled by default on all Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. However, you can explicitly enable autonegotiation to configure remote fault options manually.

NOTE:

- For Gigabit Ethernet interfaces installed in J4350 and J6350 Services Routers, when you manually configure either the link mode or speed settings, the system ignores the configuration and generates a system log message. When autonegotiation is enabled and you specify the link mode and speed, the link autonegotiates with the manually configured settings. When autonegotiation is disabled and you configure both the link mode and speed, the link operates with the manually configured settings. If you disable autonegotiation and do not manually configure the link mode and speed, the link operates at 1000 Mbps full duplex.
- When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
- On ACX Series Universal Access Routers, when the autonegotiation is disabled, the speed has to be explicitly configured to 10–100 Mbps.
- On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

Related Documentation

- Configuring Gigabit Ethernet Autonegotiation on page 416

- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Configuring Gigabit Ethernet Autonegotiation

Configuring Gigabit Ethernet Autonegotiation with Remote Fault

To configure explicit autonegotiation and remote fault, include the **auto-negotiation** statement and the **remote-fault** option at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level.

```
[edit interfaces ge-fpc/pic/port gigether-options]
(auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
  local-interface-offline>
```

Configuring Flow Control

To enable flow control, include the **flow-control** statement at the **[edit interfaces ge-*fpc/pic>/port* gigether-options]** hierarchy level. For more information, see "Configuring Flow Control" on page 45.

Configuring Autonegotiation Speed on MX Series Routers

MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs support autonegotiation of speed. The autonegotiation specified interface speed is propagated to CoS, routing protocols, and other system components. Half-duplex mode is not supported.

MX Series routers with IQ2 PICs connected to other devices require matching auto-negotiation configurations for both the PIC and for the device in order to achieve link up.

To specify the autonegotiation speed, use the **speed (auto | 1Gbps | 100Mbps | 10Mbps)** statement at the **[edit interfaces ge-*fpc/pic/port*]** hierarchy level.

To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**. If the negotiated speed and the interface speed do not match, the link will not be brought up.

If you set the autonegotiation speed **auto** option, then the port speed is negotiated.

You can disable auto MDI/MDIX using the **no-auto-mdix** statement at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level.

Use the **show interfaces ge-*fpc/pic/port* brief** command to display the auto negotiation of speed and auto MDI/MDIX states.

 Copyright © 2013, Juniper Networks, Inc.

Displaying Autonegotiation Status

To display Gigabit Ethernet interface details, including the autonegotiation status, use the operational mode command **show interfaces ge-** ***fpc/pic/port*** **extensive**.

Table 32 on page 417 and Table 33 on page 419 provide information about the autonegotiation status on local and remote routers with fiber interfaces. The status of the link and LED can vary depending on the level of autonegotiation set and the transmit and receive fiber status.

Table 32: Mode and Autonegotiation Status (Local)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Red	DOWN	
ON	ON	Default	Green	UP	No-autonegotiation
ON	OFF	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Green	UP	
ON	ON	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	

Copyright © 2013, Juniper Networks, Inc.

Table 32: Mode and Autonegotiation Status (Local) *(continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit	Green	UP	No-autonegotiation
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Green	UP	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Green	UP	Complete
OFF	ON	Explicit+RFI-Offline	Red	DOWN	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Green	UP	No-autonegotiation
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Offline	Green	UP	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	No-autonegotiation*
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Green	UP	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	

Copyright © 2013, Juniper Networks, Inc.

Table 32: Mode and Autonegotiation Status (Local) *(continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Explicit+RFI-Online	Green	UP	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	Complete

Table 33: Mode and Autonegotiation Status (Remote)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	ON	Default	Red	DOWN	
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	ON	No-autonegotiation	Red	DOWN	
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	ON	Explicit	Red	DOWN	
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	

Copyright © 2013, Juniper Networks, Inc.

Table 33: Mode and Autonegotiation Status (Remote) *(continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Offline	Red	DOWN	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	Complete
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	

Related Documentation

- Gigabit Ethernet Autonegotiation Overview on page 415
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 20

Configuring Gigabit Ethernet OTN Options

10-Gigabit Ethernet OTN Options Configuration Overview

M120, T320, T640, and T1600 routers support Optical Transport Network (OTN) interfaces, including the 10-Gigabit Ethernet DWDM OTN PIC, and provide ITU-T G.709 support. Use the **set otn-options** statement at the **[edit interfaces if-*fpc/pic/port*]** hierarchy level to configure the OTN options.

Related Documentation

- *otn-options*
- *Ethernet Interfaces*

Gigabit Ethernet OTN Options

The following example shows the configuration settings for Gigabit Ethernet OTN options:

```
[edit interfaces ge-fpc/pic/port]
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 value;
      (no-start-measurement | start-measurement;
   }
   (prbs | no-prbs);
   preemptive-fast-reroute {
      (backward-frr-enable | no-backward-frr-enable);
      (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
   }
```

Copyright © 2013, Juniper Networks, Inc.

```
    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;
  }
```

NOTE: The Gigabit Ethernet interface and the XENPAK interface support the read/write overhead bytes only for the APS/PPC (bytes 0 through 3).

You can use the following show commands to view the OTN configuration:

- **show interfaces extensive**—See the *Junos OS Operational Mode Commands* for command details.
- **show chassis hardware**—See the *Junos OS Operational Mode Commands* for command details.
- **show chassis pic**—See the *Junos OS Operational Mode Commands* for command details.

Related Documentation

- 10-Gigabit Ethernet OTN Options Configuration Overview on page 421
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 21

Configuring the Management Ethernet Interface

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 424
- Configuring the MAC Address on the Management Ethernet Interface on page 425
- Configuring MAC Filtering on PTX Series Packet Transport Routers on page 475

- *Ethernet Interfaces*

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;
    }
```


Copyright © 2013, Juniper Networks, Inc.

```
        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 423
- Configuring the MAC Address on the Management Ethernet Interface on page 425

- *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 423
- Configuring a Consistent Management IP Address on page 424
- Configuring MAC Filtering on PTX Series Packet Transport Routers on page 475

- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 22

Configuring 10-Gigabit Ethernet LAN/WAN PICs

This section contains the following topics:

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview

This section describes the main features and caveats of the 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP) and specifies which routers support this PIC.

The 10–port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) is supported on Juniper Networks T640 Core Routers, T1600 Core Routers, and T4000 Core Routers. It has the following features:

- Access to all 10-Gigabit Ethernet port counters through SNMP
- Intelligent handling of oversubscribed traffic in applications such as data centers and dense-core uplinks
- Line-rate operation for five 10-Gigabit Ethernet ports from each port group, or a total WAN bandwidth of 100 Gbps with Packet Forwarding Engine bandwidth of 50 Gbps
- Flexible encapsulation, source address and destination address media access control (MAC) filtering, source address MAC learning, MAC accounting, and MAC policing
- Interface encapsulations, such as the following:
  - **ethernet-ccc**—Ethernet cross-connect
  - **vlan-ccc**—802.1Q tagging for a cross-connect

  - **ethernet-tcc**—Ethernet translational cross-connect
  - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect
  - **extended-vlan-ccc**—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect
  - **ethernet-vpls**—Ethernet virtual private LAN service
  - **vlan-vpls**—VLAN virtual private LAN service
  - **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration

- WAN PHY features, such as the following:
  - WAN PHY mode on a per-port basis
  - Insertion and detection of path trace messages
  - Ethernet WAN Interface Sublayer (WIS) object

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).

- Single, stacked, and flexible VLAN tagging modes
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames
- Link aggregation group (LAG) on single chassis
- Interoperability with other 10-Gigabit Ethernet PICs in M Series and T Series routers in the LAN PHY and WAN PHY modes
- Interrupt-driven link-down detection mechanism
- Two-to-one oversubscription of traffic across a port group

  Traffic from 10 ingress ports to the Packet Forwarding Engine traffic is statically mapped to one of the 5 egress ports. 10 Gbps of bandwidth toward the Packet Forwarding Engine is shared by two ingress ports (called a *port group*), thereby achieving two-to-one oversubscription. This scheme provides two-to-one oversubscription across a port group and not across the entire PIC.

- Four queues per physical interface on ingress and eight queues per physical interface on egress
- A separate control queue per physical interface to ensure that the control packets are not dropped during oversubscribed traffic. The control queue can be disabled in the CLI.

 Copyright © 2013, Juniper Networks, Inc.

- Optical diagnostics
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification
- Weighted round-robin scheduling with two queue priorities (low and strict-high)
- Committed information rate and peak information rate shaping on a per-queue basis
- Excess information rate configuration for allocation of excess bandwidth
- IEEE 802.3ah Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
  - Link fault management
  - Link discovery
  - Graceful Routing Engine Switchover
- IEEE 802.3ag Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
  - Connectivity fault management (CFM)
  - Linktrace
  - Loopback
  - Graceful Routing Engine switchover (GRES)

The 10-port 10-Gigabit Ethernet LAN/WAN PIC has the following caveats:

- Source address and destination address MAC filtering takes place after oversubscription is handled.
- Oversubscription on the PIC operates across a port group of two ports and not at the PIC level.
- Queuing is not supported at the logical interface level.
- Committed information rate and peak information rate configurations are not supported at the physical interface level.
- There is limited packet buffering of 2 MB.
- Delay-bandwidth buffering configuration is not supported.
- Multifield classifiers are not supported at the PIC level.

  The multifield classification can be done at the Packet Forwarding Engine using the firewall filters, which overrides the classification done at the PIC level. The multifield classification at the Packet Forwarding Engine occurs after the PIC handles the oversubscribed traffic.
- Egress MAC policer statistics not supported.
- Byte counters are not supported at the queue level.
- Only TPID (0x8100) is supported.

Copyright © 2013, Juniper Networks, Inc.

- Line-timing mode is not supported.
- MAC-level Rx VLAN tagged frames counter is not supported.
- OAM unified in-service software upgrade (unified ISSU) is not supported.
- OAM remote loopback is not supported.

The 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) supports link aggregation. For bandwidth aggregation, load sharing, and link protection, LAG can be enabled. Once aggregated Ethernet is enabled, Link Aggregation Control Protocol (LACP) forms an aggregated bundle of member links.

Only features that are supported across all of the linked devices will be supported in the resulting LAG bundle. The following caveats apply to LAG bundles that involve 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) ports:

- Non-standard TPID for VLAN tagging is not supported, except for 0x8100.
- The number of user created IFLs is limited to 4065/PIC and 1022/port.
- Classifier tables are limited to 8 for each BA classifier type.
- Forwarding classes are limited to 8.
- The **guaranteed-rate** and **shaping-rate** statements are not supported at the IFD level.
- The **per-unit-scheduler** and **hierarchical-scheduler** statements are not supported.
- Only the **strict-high** and **low** levels of scheduling priorities are supported.
- The **excess-priority** configuration is not supported.
- The **buffer-size** configuration under **schedulers** is not supported.
- WRED is not supported.
- srTCM and trTCM are not supported.
- Shared scheduler mode is not supported.

Table 34 on page 430 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

Table 34: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs

Capability	Support
Maximum VLANs per PIC	4065
Maximum VLANs per port	1022
MAC learning per port	960
MAC accounting per port	960

Copyright © 2013, Juniper Networks, Inc.

Table 34: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs *(continued)*

Capability	Support
MAC filtering per port	960 (64 filters per physical or logical interface) 960 filters across multiple logical interfaces
MAC policers	128 ingress Mac policers 128 egress Mac policers
Classifers	Eight classifiers per PIC for each BA classifier type

Related Documentation

- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 431
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 432
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249
- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- *Ethernet Interfaces*

Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription

For 10-Gigabit Ethernet LAN/WAN PICs supporting oversubscription, oversubscribed Ethernet mode is set by default. To configure these PICs in line-rate mode, include the **linerate-mode** statement at the **[edit chassis set fpc *fpc-number* pic *pic-number*]** hierarchy level:

```
[edit chassis]
set fpc fpc-number pic pic-number linerate-mode;
```

To return to the default oversubscribed Ethernet mode, delete the **linerate-mode** statement at the **[edit chassis fpc *fpc-number* pic *pic-number*]** hierarchy level.

NOTE: When the mode of operation of a PIC is changed, the PIC is taken offline and then brought back online immediately.

The following 10-Gigabit Ethernet LAN/WAN PICs support line-rate mode:

- 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP)
- 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP)

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 427
- 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 438
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 432
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- *Ethernet Interfaces*

Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC

On a 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), a control queue is used to queue all control packets received on an ingress port. This ensures that control protocol packets do not get dropped randomly when there is congestion due to oversubscription. The following control protocols are supported:

- OSPF
- OSPF3
- VRRP
- IGMP
- RSVP
- PIM
- BGP
- BFD
- LDP
- IS-IS
- RIP
- RIPV6
- LACP
- ARP
- IPv6 NDP
- Connectivity fault management (CFM)
- Link fault management (LFM)

These control packets can either terminate locally or transit through the router. The control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. Hence, if transit control traffic is taking too much bandwidth, then it can cause drops on locally terminating control traffic, as shown in Figure 30 on page 433.

 Copyright © 2013, Juniper Networks, Inc.

Figure 30: Control Queue Rate Limiter Scenario

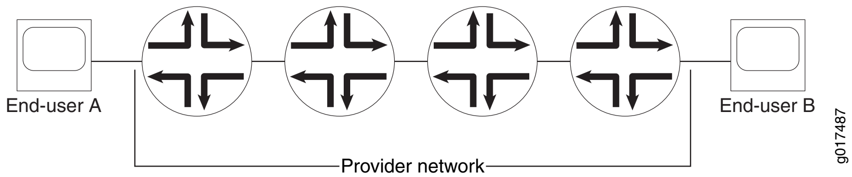

If the end users generate a mass of malicious traffic for which the port number is 179 (BGP), the router dispatches that traffic to the ingress control queue. Further, if congestion occurs in this ingress control queue due to this malicious traffic, the provider's network control packets may be affected.

In some applications, this can be perceived as a new vulnerability. To address this concern, you can disable the control queue feature. With the control queue feature disabled, you must take precautions to protect control traffic through other means, such as mapping control packets (using BA classification) to a queue that is marked strict-high or is configured with a high CIR.

You can disable the control queue for all ports on the PIC. To disable the control queue, use the **set chassis fpc *n* pic *n* no-pre-classifier** command. By default, the **no-pre-classifier** statement is not configured and the control queue is operational.

Deleting the **no-pre-classifier** statement re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

NOTE:

- This functionality is applicable both in OSE and line-rate modes.
- The control queue feature is enabled by default in both OSE and line-rate modes, which can be overridden by the user configuration.
- When the control queue is disabled, various **show queue** commands will show *control queue* in the output. However, all control queue counters are reported as zeros.
- Changing this configuration (enabling or disabling the control queue feature) results in the PIC being taken offline and brought back online.

Once the control queue is disabled, the Layer 2/Layer 3 control packets are subject to queue selection based on BA classification. However, some control protocol packets will not be classified using BA classification, because they might not have a VLAN, MPLS, or IP header. These are:

- Untagged ARP packets
- Untagged Layer 2 control packets such as LACP or Ethernet OAM
- Untagged IS-IS packets

When the control queue feature is disabled, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to the restricted queue corresponding to the forwarding class associated with queue 0, as shown in the following two examples.

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the forwarding class (FC) associated with queue 0 is "be" (based on the **forwarding-class** statement configuration). "be" maps to restricted-queue number 3 (based on the "restricted-queue" configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to ingress queue 3 (not to ingress queue 0).

```
[edit chassis]
forwarding-classes {
  queue 0 be;
  queue 1 af-low8;
  queue 2 af-high;
  queue 3 ef;
  queue 4 ops_control;
  queue 5 net_control;
  queue 6 af-low10_12;
  }
  restricted-queues {
    forwarding-class ef queue-num 0;
    forwarding-class af-low8 queue-num 1;
    forwarding-class af-low10_12 queue-num 1;
    forwarding-class af-high queue-num 2;
    forwarding-class be queue-num 3;
  }
```

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the FC associated with queue 0 is "ef" (based on the **forwarding-class** statement configuration). "ef" maps to restricted-queue number 0 (based on the **restricted-queue** statement configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets would go to ingress queue 0.

For tagged ARP, IS-IS, or Layer2 control packets, users should configure an explicit dot1p/dot1ad classifier to make sure these packets are directed to the correct queue. Without an explicit dot1p/dot1ad classifier, tagged ARP, IS-IS, or Layer 2 control packets will go to the restricted-queue corresponding to the forwarding class associated with queue 0.

```
[edit chassis]
forwarding-classes {
  queue 0 ef; <<< ef and be are interchanged
  queue 1 af-low8;
  queue 2 af-high;
  queue 3 be; <<< ef and be are interchanged
  queue 4 ops_control;
  queue 5 net_control;
  queue 6 af-low10_12;
  }
  restricted-queues {
    forwarding-class ef queue-num 0;
    forwarding-class af-low8 queue-num 1;
    forwarding-class af-low10_12 queue-num 1;
    forwarding-class af-high queue-num 2;
    forwarding-class be queue-num 3;
```

 Copyright © 2013, Juniper Networks, Inc.

}

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 427
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 431
- *no-pre-classifier*
- *Handling Oversubscription on a 10-Port 10-Gigabit Ethernet LAN/WAN PIC*
- *Ethernet Interfaces*

Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC

Table 35 on page 435 lists the scenarios of handling oversubscription on the 10-port 10-Gigabit Ethernet LAN/WAN PIC for different combinations of port groups and active ports on the PIC.

Table 35: Handling Oversubscription on 10-Gigabit Ethernet LAN/WAN PICs

Number of Port Groups with Two Active Ports (A)	Number of Port Groups with One Active Port (B)	Total Number of Ports Used on PIC (C = Ax2 + B)	Status of Oversubscription and Throughput
0	1	1	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	2	2	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	5	5	Oversubscription is not active. Each port will receive 10 Gbps throughput.
1	0	2	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
1	4	6	Oversubscription is active for the port group that has two active ports. Each port in this port group will receive 5 Gbps throughput (with default shaper configuration). For the remaining four ports, oversubscription is not active. Each port will receive 10 Gbps throughput.
3	0	6	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
5	0	10	Oversubscription is active on all 10 ports (5 port groups). Each port will receive 5 Gbps throughput (with default shaper configuration).

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 427
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 431
- Configuring Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC on page 432
- *Ethernet Interfaces*

12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC is a 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number, PF-12XGE-SFPP) on T4000 Core Routers.

The following features are supported on the 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC:

- Access to all 10-Gigabit Ethernet port counters through SNMP.
- Logical interface–level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Flexible encapsulation.
- Single, stacked, and flexible VLAN tagging modes.
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces.
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames.
- Link aggregation group (LAG) on single chassis.
- Interoperability with other 10-Gigabit Ethernet PICs on M Series and T Series routers in LAN PHY mode.
- Eight queues per physical interface on egress.
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification.
- Defining the VLAN rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.

  NOTE: **Only the Tag Protocol Identifier (TPID) 0x8100 is supported.**

- Interface encapsulations, such as the following:
  - **untagged**—Default encapsulation, when other encapsulation is not configured.
    - You can configure only one logical interface (unit 0) on the port.

 Copyright © 2013, Juniper Networks, Inc.

    - You cannot include the **vlan-id** statement in the configuration of the logical interface.

  - **vlan-tagging**—Enable VLAN tagging for all logical interfaces on the physical interface.
  - **stacked-vlan-tagging**—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
  - **ethernet-ccc**—Ethernet cross-connect.
  - **ethernet-tcc**—Ethernet translational cross-connect.
  - **vlan-ccc**—802.1Q tagging for a cross-connect.
  - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect.
  - **extended-vlan-ccc**—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect.
  - **extended-vlan-tcc**—Standard Tag Protocol Identifier (TPID) tagging for an Ethernet translational cross-connect.
  - **ethernet-vpls**—Ethernet virtual private LAN service.
  - **vlan-vpls**—VLAN virtual private LAN service.
  - **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration.

- The following Layer 3 protocols are also supported:
  - IPv4
  - IPv6
  - MPLS
- WAN PHY features, such as the following:
  - WAN PHY mode on a per-port basis.
  - Insertion and detection of path trace messages.
  - Ethernet WAN Interface Sublayer (WIS) object.

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on T4000 routers with 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+.

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC does not support:

- MAC filtering, accounting, and policing for destination MAC at the logical interface level.

  NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.
- Multiple TPIDs

Capability	Support
Maximum logical interfaces per PIC	32,000
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.
Classifers	Eight classifiers per PIC for each BA classifier type

Related Documentation

- *Ethernet Interfaces*
- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview on page 427

24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

This section describes the main features and caveats of the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP).

The following major software features are supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP):

- Twenty-four 10-Gigabit Ethernet interfaces in two-to-one oversubscription of traffic in oversubscribed mode or 12 ports in line-rate mode. For more information about oversubscribed mode and line-rate mode, see the "Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription" on page 431.
- Traffic is classified as control traffic or best-effort traffic with non-class-of-service-aware tail drops of best-effort traffic in oversubscribed mode.

  The aggregate bandwidth of all the ports together is 120 Gbps. No hard partitioning of bandwidth is done—that is, if one port group is active, it can support 120 Gbps traffic. The bandwidth for best-effort traffic is shared among all the 24 ports.

  Note that the preclassification is restricted to two traffic classes, and is not user-configurable.

 Copyright © 2013, Juniper Networks, Inc.

- All Junos OS configuration commands supported on the existing 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The output of the **show interfaces extensive** operational mode command now displays preclassification queue counters.
- Line-rate mode operation of the first 12 ports can be achieved by using the **[set chassis fpc *fpc-number* pic *pic-number* linerate-mode]** command. By default, the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ works in oversubscribed mode.
- LAN PHY mode and WAN PHY mode on a per-port basis. WAN PHY mode can be achieved by using the **[set interfaces *interface-name* framing wan-phy]** command.
- WAN PHY features, such as the following:
  - Insertion and detection of path trace messages.
  - Ethernet WAN Interface Sublayer (WIS) object.
- Aggregated Ethernet is supported only in line-rate mode.
- Link aggregation group (LAG) is supported only in line-rate mode.
- 4000 logical interfaces per physical interface and 32,000 logical interfaces per chassis.
- Access to all 10-Gigabit Ethernet port counters through SNMP.

NOTE: **Graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are now supported on T4000 routers.**

Related Documentation

- *Ethernet Interfaces*
- 12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview on page 436
- Configuring Line-Rate Mode on 10-Gigabit Ethernet LAN/WAN PICs Supporting Oversubscription on page 431

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 23

Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength

Ethernet DWDM Interface Wavelength Overview

MX960, M320, M120, T320, and T640 routers using a 10-Gigabit Ethernet DWDM PIC and PTX Series routers using a 100-Gigabit Ethernet DWDM OTN PIC support DWDM interfaces. You can configure DWDM interfaces with full C-band International Telecommunication Union (ITU)-Grid tunable optics, as defined in the following specifications:

- *Intel TXN13600 Optical Transceiver I2C Interface and Customer EEPROM Preliminary Specification*, July 2004.
- *I2C Reference Document for 300-Pin MSA 10G and 40G Transponder*, Edition 4, August 04, 2003.

By default, the wavelength is 1550.12 nanometers (nm), which corresponds to 193.40 terahertz (THz).

Related Documentation

- Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength on page 441

- *Ethernet Interfaces*

Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength

To configure the wavelength on a 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interface, include the **wavelength** statement at the **[edit interfaces *media-type-fpc/pic/port* optics-options]** hierarchy level:

```
[edit interfaces ge-0/0/0 optics-options]
wavelength nm;
```

Copyright © 2013, Juniper Networks, Inc.

For interface diagnostics, you can issue the **show interfaces diagnostics optics** ***media-type-fpc/pic /port*** operational mode command.

Table 36 on page 442 shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

Table 36: Wavelength-to-Frequency Conversion Matrix

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
1528.77	196.10	1540.56	194.60	1552.52	193.10
1529.55	196.00	1541.35	194.50	1553.33	193.00
1530.33	195.90	1542.14	194.40	1554.13	192.90
1531.12	195.80	1542.94	194.30	1554.94	192.80
1531.90	195.70	1543.73	194.20	1555.75	192.70
1532.68	195.60	1544.53	194.10	1556.56	192.60
1533.47	195.50	1545.32	194.00	1557.36	192.50
1534.25	195.40	1546.12	193.90	1558.17	192.40
1535.04	195.30	1546.92	193.80	1558.98	192.30
1535.82	195.20	1547.72	193.70	1559.79	192.20
1536.61	195.10	1548.52	193.60	1560.61	192.10
1537.40	195.00	1549.32	193.50	1561.42	192.00
1538.19	194.90	1550.12	193.40	1562.23	191.90
1538.98	194.80	1550.92	193.30	1563.05	191.80
1539.77	194.70	1551.72	193.20	1563.86	191.70

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 441
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 24

Configuring 10-Gigabit Ethernet Framing

10-Gigabit Ethernet Framing Overview

The 10-Gigabit Ethernet interfaces support operation in two modes:

- 10GBASE-R, LAN Physical Layer Device (LAN PHY)
- 10GBASE-W, WAN Physical Layer Device (WAN PHY)

When the external interface is running in LAN PHY mode, it bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface. When the external interface is running in WAN PHY mode, it uses the WIS sublayer to transport 10-Gigabit Ethernet frames in an OC192c SONET payload.

WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600, T4000 and PTX Series Packet Transport routers only.

NOTE: **The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).**

Although the external interface provides a lower throughput when running in WAN PHY mode because of the extra SONET overhead, it can interoperate with SONET section or line level repeaters. This creates an advantage when the interface is used for long-distance, point-to-point 10-Gigabit Ethernet links. When the external interface is running in WAN PHY mode, some SONET options are supported. For information about SONET options supported on this interface, see *Configuring SONET Options for 10-Gigabit Ethernet Interfaces*.

NOTE: SONET or SDH framing mode configuration **framing (sdh | sonet)** is not applicable on the 10-Gigabit Ethernet ports. Configuring the **wan-phy** framing mode on the 10-Gigabit Ethernet ports allows the interface to accept SONET or SDH frames without further configuration.

Related Documentation

- *Configuring SONET/SDH Framing Mode*
- Configuring 10-Gigabit Ethernet Framing on page 444
- Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces on page 445
- *Ethernet Interfaces*

Configuring 10-Gigabit Ethernet Framing

The 10-Gigabit Ethernet interfaces uses the interface type **xe-*fpc/pic/port***. On single port devices, the port number is always zero.

The **xe-*fpc/pic/port*** interface inherits all the configuration commands that are used for gigabit Ethernet (**ge-*fpc/pic/port***) interfaces.

To configure LAN PHY or WAN PHY operating mode, include the **framing** statement with the **lan-phy** or **wan-phy** option at the **[edit interfaces xe-*fpc /pic/*0]** hierarchy level.

```
[edit interfaces xe-fpc/pic/0 framing]
framing (lan-phy | wan-phy);
```

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- On PTX Series Transport Routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.

To display interface information, use the operational mode command **show interfaces xe-*fpc/pic/port* extensive**.

 Copyright © 2013, Juniper Networks, Inc.

NOTE:

- SONET or SDH framing mode configuration **framing (sdh | sonet)** is not applicable on the 10-Gigabit Ethernet ports. Configuring the **wan-phy** framing mode on the 10-Gigabit Ethernet ports allows the interface to accept SONET or SDH frames without further configuration.
- If you configure the WAN PHY mode on an aggregated Ethernet interface, you must set the aggregated Ethernet link speed to OC192.

Related Documentation

- **framing on page 576**
- 10-Gigabit Ethernet Framing Overview on page 443
- Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces on page 445
- *Ethernet Interfaces*

Understanding WAN Framing for 10-Gigabit Ethernet Trio Interfaces

If you use the **wan-phy** statement option at the **[edit interfaces xe-*fpc/pic*/0 framing]** hierarchy level to configure Trio WAN mode framing for 10-Gigabit Ethernet interfaces, then the alarm behavior of the link, although in full compliance with the IEEE 802.3ae 10-Gigabit Ethernet standard, might not be as expected.

In particular:

- The interface does not distinguish between loss of light (LOL), loss of phase lock loop (PLL), or loss of signal (LOS). If a loss of PLL or LOS alarm occurs, then both PLL and LOS alarms are raised. LOL is also raised because there is no separate LOL indication from the hardware.
- The interface does not raise LOS, PLL, or LOL alarms when the fiber in disconnected from the interface port. You must remove the hardware to raise this alarm.
- The interface line-level alarm indicator signal (AIS-L) is not always raised in response to a loss of framing (LOF) defect alarm.
- If the AIS-L or path-level AIS (AIS-P) occurs, the interface path-level loss of code delineation (LCD-P) is not detected. LCD-P is seen during the path-level remote defect indicator (RDI-P) alarm.
- If an AIS-L alarm occurs, the AIS-P is not detected, but the LOP alarm is detected.

None of the alarm issues are misleading, but they make troubleshooting the root cause of problems more complex.

Related Documentation

- **framing on page 576**
- Configuring 10-Gigabit Ethernet Framing on page 444
- 10-Gigabit Ethernet Framing Overview on page 443
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 25

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

10-Gigabit Ethernet Notification of Link Down Alarm Overview

Notification of link down alarm generation and transfer is supported for all 10-Gigabit Ethernet PIC interfaces in M120, M320. and T Series routers.

Related Documentation

- Configuring 10-Gigabit Ethernet Notification of Link Down Alarm on page 447
- *Ethernet Interfaces*

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

To configure this option, include the **asynchronous-notification** statement at the **[edit interfaces ge-** ***fpc/pic/port*** **gigether-options]** hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
  gigether-options {
    asynchronous-notification;
  }
}
```

Related Documentation

- 10-Gigabit Ethernet Notification of Link Down Alarm Overview on page 447
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 26

Configuring 10-Gigabit Ethernet Notification of Link Down for Optics Alarms

10-Gigabit Ethernet Notification of Link Down for Optics Options Overview

Notification of link down is supported for IQ2 10-Gigabit Ethernet interfaces and MX Series DPCs. You can use link down notification to help identify optical link connectivity problems.

For information on configuring link down notification, see "Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning" on page 449.

Related Documentation

- Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning on page 449
- *Ethernet Interfaces*

Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning

To configure this option, include the **alarm** or **warning** statement at the **[edit interfaces ge-** ***fpc/pic/port*** **optics-options]** hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
  optics-options {
    alarm alarm-name {
      (syslog | link-down);
      }
    warning warning-name {
      (syslog | link-down);
      }
  }
```

```
}
```

Related Documentation

- *alarm*
- *warning*
- 10-Gigabit Ethernet Notification of Link Down for Optics Options Overview on page 449
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

CHAPTER 27

Configuring 100-Gigabit Ethernet PICs/MICs

This section contains the following topics:

100-Gigabit Ethernet Type 4 PIC with CFP Overview

The 100-Gigabit Ethernet PIC is a 1-port 100-Gigabit Ethernet Type 4 PIC with 100-gigabit small form-factor pluggable transceiver (CFP) with model number PD-1CE-CFP-FPC4. The 100-Gigabit Ethernet PIC occupies FPC slots 0 and 1 in the T1600-FPC4-ES FPC. This PIC is available only as packaged in an assembly with the T1600-FPC4-ES FPC. For information on supported transceivers and hardware, see *100-Gigabit Ethernet PIC with CFP (T1600 Router)*.

The 100-Gigabit Ethernet PIC supports flexible encapsulation and MAC accounting.

MAC learning, MAC policing, and Layer 2 rewrite functionality are not supported.

The ingress flow can be filtered based on the VLAN source and destination addresses. Ingress frames can also be classified according to VLAN, stacked VLAN, source address, VLAN source address, and stacked VLAN source address. VLAN manipulation on egress frames are supported on both outer and inner VLAN tags.

The following features are supported:

- The following encapsulation protocols are supported:

- Layer 2 protocols
  - Ethernet CCC, Ethernet TCC, Ethernet VPLS
  - VLAN CCC
  - Extended VLAN TCC
  - VLAN VPLS
  - Flexible Ethernet service
- Layer 3 protocols
  - IPv4
  - Ipv6
  - MPLS

- CFP MSA compliant MDIO control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) is supported in all PIC and chassis configurations.
- Interface creation:
  - When the PIC, is brought online, the router creates two 50 gigabit capable interfaces, **et-*x*/0/0:0** and **et-*x*/0/0:1**, where *x* represents the FPC slot number. Each physical interface represents two internal 50 gigabit Ethernet Packet Forwarding Engines. Two logical interfaces are configured under each physical interface.
  - Packet Forwarding Engine 0 is physical interface 0, Packet Forwarding Engine 1 is physical interface 1
- 802.3 link aggregation:

  Same rate or same mode link aggregation:

  - Two logical interfaces are created for each 100-Gigabit Ethernet PIC. To utilize bandwidth beyond 50 gigabits per second, an aggregate interface must be explicitly configured on the 100-Gigabit Ethernet PIC that includes the two 50 gigabit interfaces.
  - Each 100 gigabit Ethernet aggregate consumes one of the router-wide aggregated Ethernet device pools. The number of 100-Gigabit Ethernet PICs cannot exceed the router-wide limit, which is 128 for Ethernet.
  - In each aggregate bundle, each 100-Gigabit Ethernet PIC consumes two members. Hence, an aggregate bundle that consists purely of 100-Gigabit Ethernet PICs supports a maximum of half of the software limit for the number of members. Therefore, with a maximum of 16 links, up to 8 100-Gigabit Ethernet links are supported.
  - Combining 100-Gigabit Ethernet PICs into aggregate interfaces with other Ethernet PICs is not permitted. However, other Ethernet PICs can also be configured within

 Copyright © 2013, Juniper Networks, Inc.

the same T1600 with 100-Gigabit Ethernet PICs, and used in separate aggregate interfaces.

- Multiple (Juniper Networks) Type 4 100-Gigabit Ethernet PICs on a T1600 router can be combined into a static aggregated Ethernet bundle to connect to a different type of 100 gigabit Ethernet PIC on a remote router (Juniper Networks or other vendors). LACP is not supported in this configuration.

Mixed rate or mixed mode link aggregation:

- Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC.
- Static link protection and Link Aggregation Control Protocol (LACP) is supported on mixed aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC.
- When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC, ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC.

- Software Packet Forwarding Engine—Supports all Gigabit Ethernet PIC classification, firewall filter, queuing model, and rewrite functionality.
- Egress traffic performance—Maximum egress throughput is 100 gigabits per second on the physical interface, with 50 gigabits per second on the two assigned logical interfaces.
- Ingress traffic performance—Maximum ingress throughput is 100 gigabits per second on the physical interface, with 50 gigabits per second on the two assigned logical interfaces. To achieve 100 gigabits per second ingress traffic performance, use one of the interoperability modes described below. For example, if VLAN steering mode is not used when connecting to a remote 100 gigabits per second interface (that is on a different 100 gigabits per second PIC on a Juniper Networks router or a different vendor's equipment), then all ingress traffic will try to use one of the 50 gigabits per second Packet Forwarding Engines, rather than be distributed among the two 50 gigabits per second Packet Forwarding Engines, resulting in a total of 50 gigabits per second ingress performance.
- Interoperability modes—The 100-Gigabit Ethernet PIC supports interoperability with through configuration in one of the following two forwarding option modes:
  - *SA multicast mode*—In this mode, the 100-Gigabit Ethernet PIC supports interconnection with other Juniper Networks 100-Gigabit Ethernet PICs (Model: PD-1CE-CFP) interfaces only.
  - *VLAN steering mode*—In this mode, the 100-Gigabit Ethernet Type 4 PIC with CFP supports interoperability with 100 gigabit Ethernet interfaces from other vendors only.

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP on page 456
- *T1600 Core Router*
- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

MPC3E MIC Overview

The MPC3E supports two separate slots for MICs. MICs provide the physical interface and are installed into the MPCs.

The MPC3E supports these MICs as field replaceable units (FRUs):

- *100-Gigabit Ethernet MIC with CFP* (model number MIC3-3D-1X100GE-CFP)
- *100-Gigabit Ethernet MIC with CXP* (model number MIC3-3D-1X100GE-CXP)
- 10-port 10-Gigabit Ethernet MIC with SFPP (model number MIC3-3D-10XGE-SFPP)
- 2-port *40-Gigabit Ethernet MIC with QSFP+* (model number MIC3-3D-2X40GE-QSFPP)

The MPC3E has two separate configurable MIC slots. Each MIC corresponds to a single PIC and the mapping between the MIC and PIC is 1 to 1 (one MIC is treated as one PIC). The MIC plugged into slot 0 corresponds to PIC 0 and the MIC plugged into slot 1 corresponds to PIC 2.

The MPC3E also supports these legacy MICs:

- 20-port *Gigabit Ethernet MIC with SFP* (model number MIC-3D-20GE-SFP)
- 2-port *10-Gigabit Ethernet MICs with XFP* (model number MIC-3D-2XGE-XFP)

The 100-Gigabit Ethernet CFP MIC supports the IEEE standards—compliant 100BASE-LR4 interface, using the 100G CFP optical transceiver modules for connectivity. The 100-Gigabit Ethernet CXP MIC supports the 100BASE-SR10 interface, using 100-Gigabit CXP optical transceiver modules for connectivity. The 2-port 40-Gigabit Ethernet QSFPP MIC supports the 40BASE-SR4 interface and uses quad small form-factor pluggable (QSFPP) optical transceivers for connectivity. The 10-port 10-Gigabit Ethernet SFPP MIC uses SFP+ optical transceiver modules for connectivity.

For detailed information about each MIC, see *100-Gigabit Ethernet MIC with CFP*, *100-Gigabit Ethernet MIC with CXP*, *40-Gigabit Ethernet MIC with QSFP+*. For information about supported hardware and transceivers, see *MPC3E*.

The MPC3E supports these features:

- Optical diagnostics and related alarms
- Virtual Router Redundancy Protocol (VRRP) support
- IEEE 802.1Q virtual LANs (VLANs) support
- Synchronous Ethernet

 Copyright © 2013, Juniper Networks, Inc.

- Remote monitoring (RMON) and Ethernet statistics (EtherStats)
- Source MAC learning
- MAC accounting and policing—Dynamic local address learning of source MAC addresses
- Flexible Ethernet encapsulation
- Multiple Tag Protocol Identifiers (TPIDs)

NOTE: The MPC3E supports Ethernet interfaces only. SONET interfaces are not supported.

For information about the supported and unsupported Junos OS features for this MPC, see "Protocols and Applications Supported by the MPC3E (MX-MPC3E)" in the *MX Series Interface Module Reference*.

Related Documentation

- *MPC3E on MX Series Routers Overview*
- *Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC3E*
- *100-Gigabit Ethernet MIC with CFP*
- *100-Gigabit Ethernet MIC with CXP*
- 2-port*40-Gigabit Ethernet MIC with QSFP+*
- 2-port *10-Gigabit Ethernet MICs with XFP*
- *MX Series Interface Module Reference*

Copyright © 2013, Juniper Networks, Inc.

Configuring 100-Gigabit Ethernet Type 4 PIC With CFP

You can configure the following features on the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4):

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting in RX
- Channels defined by two stacked VLAN tags
- Channels defined by flex-vlan-tagging
- IP service for stacked VLAN tags
- Layer 2 rewrite

The following features are not supported on the 100-Gigabit Ethernet Type 4 PIC with CFP:

- Multiple TPID
- IP service for non-standard TPID
- MAC learning
- MAC policing

NOTE:

- For the 100-Gigabit Ethernet Type 4 PIC with CFP, only the PIC0 **online** and **offline** CLI commands are supported. The PIC1 **online** and **offline** CLI commands are not supported.
- Each 100-Gigabit Ethernet Type 4 PIC with CFP creates two et- physical interfaces, defined as 50-gigabit physical interfaces in the Routing Engine and Packet Forwarding Engine. By default, these are independent physical interfaces and are not configured as an aggregated Ethernet interface.

 Copyright © 2013, Juniper Networks, Inc.

To configure a 100-Gigabit Ethernet Type 4 PIC with CFP:

1. Perform the media configuration:

   The 100-Gigabit Ethernet Type 4 PIC with CFP features a 100 gigabit per second pipe. The media-related configuration commands for **et-*x*/0/0:0** and **et-*x*/0/0:1** must both be configured at the same time and configured with the same value, otherwise the commit operation fails.

   When configuring to activate or deactivate the interface, if the interface contains the described media-related configuration, it must activate and deactivate both units 0 and 1 at the same time, otherwise the commit operation fails.

   The following media configuration commands have the above described restriction:

   - # **set interfaces et-x/0/0:1 disable**
   - # **set interfaces et-x/0/0:1 gigether-options loopback**
   - # **set interfaces et-x/0/0:1 mtu yyy**

   Due to an MTU restriction, the vlan-tagging and flexible-vlan-tagging configuration on et-x/0/0:0 and et-x/0/0:1 must be same, otherwise the commit operation fails.

2. Specify the logical interfaces:

   a. Two physical interfaces are created when the 100-Gigabit Ethernet Type 4 PIC with CFP is brought online (**et-*x*/0/0:0** and **et-*x*/0/0:1**, where *x* represents the FPC slot number). Each physical interface represents two internal 50-gigabit Ethernet Packet Forwarding Engines.

   b. Two logical interfaces are configured under each physical interface: Packet Forwarding Engine 0 is physical interface 0 and Packet Forwarding Engine 1 is physical interface 1.

3. Configure the 802.3 link aggregation:

   a. The 100-Gigabit Ethernet PIC supports aggregated Ethernet configuration to achieve higher throughput capability, whereby configuration is similar to the 1G/10G aggregated Ethernet interface configuration.

   b. Two physical interfaces are created for each 100-Gigabit Ethernet Type 4 PIC with CFP. To utilize bandwidth beyond 50 gigabits, a same rate and same mode aggregated Ethernet interface must be explicitly configured on the 100-Gigabit Ethernet Type 4 PIC with CFP that includes these two 50-gigabit interfaces.

   c. Each 100-Gigabit Ethernet Type 4 PIC with CFP aggregate consumes one of the router-wide aggregated Ethernet device pools. In Junos OS with 100-Gigabit Ethernet PICs, you cannot exceed the router limit of 128 Ethernet PICs.

   d. In each aggregated bundle, each 100-Gigabit Ethernet Type 4 PIC with CFP consumes two aggregate members. Hence, an aggregated bundle consisting of only one 100-Gigabit Ethernet Type 4 PIC with CFP supports only up to half of the Junos OS limit for the number of members. The Junos OS supports a maximum of 16 links for up to 8 100-Gigabit Ethernet Type 4 PIC with CFP links.

Copyright © 2013, Juniper Networks, Inc.

NOTE:

The 100-Gigabit Ethernet Type 4 PIC with CFP has the following restrictions for same rate and same mode aggregated Ethernet configuration:

- Both physical interfaces belonging to the same 100-Gigabit Ethernet PIC must be included in the same aggregated Ethernet physical interfaces. The aggregation of the 100-Gigabit Ethernet PIC interface is always an even number of physical interfaces.
- The 100-Gigabit Ethernet PIC physical interface cannot be configured in the aggregated interface with any other type of physical interface.
- The maximum supported number of aggregated 100-Gigabit Ethernet PIC interfaces is half of the number that the Junos OS supports for 1G/10G aggregated Ethernet. For example, if Junos OS supports 16 ports of 10-gigabit Ethernet aggregation, it supports 8 ports of 100-Gigabit Ethernet PIC aggregation. This is because each port of the 100-Gigabit Ethernet PIC port using 2 physical interfaces (et-x/0/0:0 and et-x/0/0:1), where each physical interface represents 50 gigabits of traffic capacity.

e. Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC. When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC, ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.

 Copyright © 2013, Juniper Networks, Inc.

NOTE:

The 100-Gigabit Ethernet Type 4 PIC with CFP has the following restrictions for mixed rate and mixed mode aggregated Ethernet configuration:

- A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
- Traffic distribution is based on the hash calculated on the egress packet header. Hash range is fairly distributed according to member links' speed. This guarantees hash fairness but it does not guarantee fair traffic distribution depending on the rate of the egress streams.
- Packets are dropped when the total throughput of the hash flow exiting a member link (or multiple hash flows exiting a single member link) exceeds the link speed of the member link. This can happen when egress member link changes because of a link failure and the hash flow switches to a member link of speed that is less than the total throughput of the hash flow.
- Rate-based CoS components such as scheduler, shaper, and policer are not supported on mixed rate aggregated Ethernet links. However, the default CoS settings are supported by default on the mixed rate aggregated Ethernet links.
- Load balancing is performed at the ingress Packet Forwarding Engine. Therefore, you must ensure that the egress traffic on the aggregated Ethernet link enters through the hardware platforms that support mixed aggregated Ethernet bundles.
- Mixed aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed aggregated Ethernet load balancing is configured at egress.
- Load balancing of the egress traffic across the member links of a mixed rate aggregated Ethernet link is proportional to the rates of the member links.
- Egress multicast load balancing is not supported on mixed aggregated Ethernet interfaces.
- Changing the **edit interfaces ae*x* aggregated-ether-options link-speed** configuration of a mixed aggregated Ethernet link, which is configured on the supported interfaces of on T640, T1600, T4000, and TX Matrix Plus routers, leads to aggregated Ethernet link flapping.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC, changing aggregated Ethernet link protection configurations leads to aggregated Ethernet link flapping.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP.

- The **show interfaces ae*x*** command displays the link speed of the aggregated Ethernet interface, which is the sum of the link speeds of all the active member links.

4. Configure the Packet Forwarding Engine features:

   a. The 100-Gigabit Ethernet Type 4 PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see "Configuring Gigabit Ethernet Policers" on page 397, "Configuring MAC Address Filtering" on page 401, and "Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview" on page 197.

NOTE: When using the **show interfaces extensive** command with a 100-Gigabit Ethernet Type 4 PIC with CFP, the "Filter statistics" section will not be displayed because the hardware does not include those counters.

Related Documentation

- 100-Gigabit Ethernet Type 4 PIC with CFP Overview on page 451
- Configuring Gigabit Ethernet Policers on page 397
- Configuring MAC Address Filtering on page 401
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 197

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

In Junos OS Release 10.4 and later, you can configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with routers using 100 gigabit Ethernet interfaces from other vendors by using the **forwarding-mode** statement with the **vlan-steering** option at the **[edit chassis fpc *slot* pic *slot*]** hierarchy level. On ingress, the router compares the outer VLAN ID against the user-defined VLAN ID and VLAN mask combination and steers the packet accordingly. You can program a custom VLAN ID and corresponding mask for PFE0.

General information on the VLAN steering mode:

- In VLAN steering mode, the SA multicast parameters are not used for packet steering.
- In SA multicast bit steering mode, the VLAN ID and VLAN masks are not used for packet steering.
- Configuration to set the packet distribution mode and VLAN steering rule is done through CLI commands. Both CLI commands result in a PIC reboot.
- There are three possible tag types of ingress packet:
  - Untagged ingress packet—The packet is sent to PFE1.
  - Ingress packet with one VLAN—The packet is forwarded to the corresponding PFE based on the VLAN ID.

 Copyright © 2013, Juniper Networks, Inc.

    - Ingress packet with two VLANs—The packet is forwarded to the corresponding PFE based on the outer VLAN ID.

- If no VLAN rule is configured, all tagged packets are distributed to PFE0.
- VLAN rules describe how the router distributes packets. Two VLAN rules are provided by the CLI:
    - Odd-Even rule—Odd number VLAN IDs go to PFE1; even number of VLAN IDs go to PFE0.
    - Hi-Low rule—VLAN IDs 1 through 2047 go to PFE0; VLAN IDs 2048 through 4096 go to PFE1.
- When the 100-Gigabit Ethernet Type 4 PIC with CFP is configured in VLAN steering mode, it can be configured in a two physical interfaces mode or in aggregate Ethernet (AE) mode:
    - Two physical interfaces mode—When the PIC is in the two physical interfaces mode, it creates the physical interfaces **et-*x/0/0*:*0*** and **et-*x/0/0*:*1***. Each physical interface can configure its own logical interface and VLAN. The CLI enforces the following restrictions at the commit time:
        - The VLAN ID configuration must comply with the selected VLAN rule.
        - The previous restriction implies that the same VLAN ID cannot be configured on both physical interfaces.
    - AE mode—When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one AE physical interface. The PIC egress traffic is based on an AE internal hash algorithm. The PIC ingress traffic steering is based on the customized VLAN ID rule. The CLI enforces the following restrictions at the commit time:
        - The PICs AE working in VLAN steering mode includes both links of that PIC, and only the links of that PIC.
        - The PIC AE working in SA multicast steering mode can include more than one 100-Gigabit Ethernet Type 4 PIC with CFP to achieve more than 100 gigabit Ethernet capacity.

To configure SA multicast mode, use the **set chassis fpc *slot* pic *slot* forwarding-mode sa-multicast** command.

SA Multicast Mode

To configure SA multicast mode on a Juniper Networks 100-Gigabit Ethernet Type 4 PIC with CFP in FPC 0, PIC 0 for interconnection with another Juniper Networks 100-Gigabit Ethernet PIC, use the **set chassis fpc *slot* pic *slot* forwarding-mode sa-multicast** command. You can use the **show forwarding-mode** command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
forwarding-mode {
  sa-multicast;
}
```

Copyright © 2013, Juniper Networks, Inc.

VLAN Steering Mode

To configure the Juniper Networks 100-Gigabit Ethernet Type 4 PIC with CFP for VLAN steering mode for interoperation with a 100 gigabit Ethernet interface from another vendor's router, use the **set chassis fpc *slot* pic *slot* forwarding-mode vlan-steering** command with the **vlan-rule (high-low | odd-even)** statement. You can use the **show forwarding-mode** command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
forwarding-mode {
  vlan-steering {
    vlan-rule odd-even;
  }
}
```

Related Documentation

- forwarding-mode (100-Gigabit Ethernet) on page 575
- sa-multicast (100-Gigabit Ethernet) on page 640
- vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 683
- vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 684

100-Gigabit Ethernet Type 5 PIC with CFP Overview

The 100-Gigabit Ethernet PIC is a 1-port 100-Gigabit Ethernet Type 5 PIC with C form-factor pluggable transceiver (CFP) with model number PF-1CGE-CFP.

The following features are supported on 100-Gigabit Ethernet Type 5 PIC with CFP:

- Access to all 100-Gigabit Ethernet port counters through SNMP.
- Logical interface–level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Channels defined by two stacked VLAN tags.
- Channels defined by **flex-vlan-tagging**.
- IP service for stacked VLAN tags.
- Defining the rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.

NOTE: Only the Tag Protocol Identifier (TPID) 0x8100 is supported.

- Interface encapsulations, such as the following:
  - **untagged**—Default encapsulation, when other encapsulation is not configured.
    - You can configure only one logical interface (unit 0) on the port.
    - You cannot include the **vlan-id** statement in the configuration of the logical interface.
  - **vlan-tagging**—Enable VLAN tagging for all logical interfaces on the physical interface.

 Copyright © 2013, Juniper Networks, Inc.

  - **stacked-vlan-tagging**—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
  - **ethernet-ccc**—Ethernet cross-connect.
  - **ethernet-tcc**—Ethernet translational cross-connect.
  - **vlan-ccc**—802.1Q tagging for a cross-connect.
  - **vlan-tcc**—Virtual LAN (VLAN) translational cross-connect.
  - **extended-vlan-ccc**—Standard TPID tagging for an Ethernet cross-connect.
  - **extended-vlan-tcc**—Standard TPID tagging for an Ethernet translational cross-connect.
  - **flexible-ethernet-services**—Allows per-unit Ethernet encapsulation configuration.
  - **ethernet-vpls**—Ethernet virtual private LAN service.
  - **vlan-vpls**—VLAN virtual private LAN service.
- The following Layer 3 protocols are also supported:
  - IPv4
  - IPv6
  - MPLS
- CFP Multi-Source Agreement (MSA) compliant Management Data Input/Output (MDIO) control features (transceiver dependent).
- 802.3 link aggregation:
  - The configuration of the 100-Gigabit Ethernet Type 5 PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
- Interoperability mode—Interoperability with the 100-Gigabit Ethernet Type 4 PIC with CFP through configuration in **sa-multicast** forwarding mode.
- Juniper Networks enterprise-specific Ethernet Media Access Control (MAC) MIB
- The 100-Gigabit Ethernet Type 5 PIC with CFP supports all Gigabit Ethernet PIC classification, firewall filters, queuing model, and Layer 2 rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see “Configuring Gigabit Ethernet Policers” on page 397, “Configuring MAC Address Filtering” on page 401, and “Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview” on page 197.
- A Type 5 FPC can support up to two 100-Gigabit Ethernet PICs. Both the PICs (that is, PIC 0 and PIC 1) can be offline or online independently.

Copyright © 2013, Juniper Networks, Inc.

The following features are not supported on the 100-Gigabit Ethernet Type 5 PIC with CFP:

- MAC filtering, accounting, and policing for destination MAC at the logical interface level.

NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This configuration enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.
- Multiple TPIDs.
- IP service for nonstandard TPID.

Table 37 on page 464 lists the capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP.

Table 37: Capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP

Capability	Support
Maximum logical interfaces per PIC	4093
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.

Related Documentation

- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP on page 456
- Configuring Gigabit Ethernet Policers on page 397
- Configuring MAC Address Filtering on page 401
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 197

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP by:

- Enabling source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP.
- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the 100-Gigabit Ethernet

 Copyright © 2013, Juniper Networks, Inc.

PIC PF-1CGE-CFP. The 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 looks at the bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. When the PIC sends out a packet, the multicast bit is set based on the egress Packet Forwarding Engine number (0 or 1).

The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.

PD-1CE-CFP-FPC4 uses two 50 Gpbs Packet Forwarding Engines to achieve 100 Gbps throughput. The 50-Gigabit Ethernet physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both the physical interfaces. You must configure the physical interfaces on PD-1CE-CFP-FPC4 in static link aggregation group (LAG) mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to the 100-Gigabit Ethernet PIC PF-1CGE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

NOTE: If you try to enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP without configuring PD-1CE-CFP-FPC4 (with two 50-Gigabit Ethernet interfaces) in static LAG mode, then there are issues in forwarding or routing protocols. For example, if you create two untagged logical interfaces—one each on the two 50-Gigabit Ethernet interfaces—on PD-1CE-CFP-FPC4 and one untagged logical interface on PF-1CGE-CFP, then PF-1CGE-CFP does not learn about one of the 50-Gigabit Ethernet interfaces on PD-1CE-CFP-FPC4.

Related Documentation

- forwarding-mode on page 575
- sa-multicast on page 640
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 465

- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP by performing the following tasks:

- Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP on page 466
- Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface on page 466

Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP, you need to enable source address (SA) multicast bit steering mode on PF-1CGE-CFP.

To configure SA multicast mode on PF-1CGE-CFP:

1. Specify the FPC and PIC information on the chassis.

   [edit]
   user@host# **edit chassis fpc *slot* pic *slot***

   For example:

   [edit]
   user@host# **edit chassis fpc 1 pic 0**

2. Configure the interoperation mode (SA multicast bit steering mode).

   [edit chassis fpc *slot* pic *slot*]
   user@host# **set forwarding-mode sa-multicast**

   For example:

   [edit fpc 1 pic 0]
   user@host# **set forwarding-mode sa-multicast**

3. Verify the configuration.

```
[edit ]
user@host# show chassis
  fpc 1 {
    pic 0 {
      forwarding-mode {
        sa-multicast;
      }
    }
  }
```

NOTE: **The default packet steering mode for the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.**

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP or P1-PTX-2-100GE-CFP, you need to configure the two 50-Gigabit Ethernet physical interfaces on PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface. This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to PF-1CGE-CFP or P1-PTX-2-100GE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

Copyright © 2013, Juniper Networks, Inc.

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-*fpc*/*pic*/0:0 and et-*fpc*/*pic*/0:1, where *fpc* is the FPC slot number and *pic* is the PIC slot number. For example, to configure two physical interfaces for PIC slot 0 in FPC slot 5:

1. Specify the number of aggregated Ethernet interfaces to be created.

   ```
   [edit chassis]
   user@host# set aggregated devices ethernet device-count count
   ```

   For example:

   ```
   [edit chassis]
   user@host# set aggregated devices ethernet device-count 1
   ```

2. Specify the members to be included within the aggregated Ethernet bundle.

   ```
   [edit interfaces ]
   user@host# set interface-name gigether-options 802.3ad bundle
   ```

   The following example shows how to configure two physical interfaces for PIC 0 on a T1600 router.

   ```
   [edit interfaces ]
   user@host# set et-5/0/0:0 gigether-options 802.3ad ae0
   user@host# set et-5/0/0:1 gigether-options 802.3ad ae0
   ```

3. Verify the configuration at the chassis.

   ```
   [edit ]
   user@host# show chassis
     aggregated-devices {
       ethernet {
         device-count 1;
       }
     }
   ```

4. Verify the configuration at the interface.

   ```
   [edit ]
   user@host# show interfaces
     et-5/0/0:0 {
       gigether-options {
         802.3ad ae0;
       }
     }
     et-5/0/0:1 {
       gigether-options {
         802.3ad ae0;
       }
     }
   ```

Related Documentation

• Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 464

- *100-Gigabit Ethernet PIC with CFP (T1600 Router)*
- *100-Gigabit Ethernet PIC with CFP (T4000 Router)*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 28

Configuring 40-Gigabit Ethernet PICs

This section contains the following topics:

40-Gigabit Ethernet PIC Overview

The 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP) is a 1-port 40-Gigabit Ethernet Type 4 PIC with C form-factor pluggable transceiver (CFP) optics supported on T640, T1600, and T4000 routers. The 40-Gigabit Ethernet PIC occupies FPC slot 0 or 1 in the Type 4 FPC and it is similar to any regular PIC such as the 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP) PIC. The CFP information appears under the PIC information in the show command output.

The 40-Gigabit Ethernet PIC with CFP supports flexible Ethernet services encapsulation and MAC accounting.

MAC learning, MAC policing, and Layer 2 rewrite features are not supported.

The 40-Gigabit Ethernet PIC with CFP supports the following features:

- Encapsulation protocols such as:
  - Layer 2 protocols
    - Ethernet CCC, Ethernet TCC, and Ethernet VPLS
    - VLAN CCC
    - Extended VLAN TCC
    - VLAN VPLS
    - Flexible Ethernet service
  - Layer 3 protocols
    - IPv4
    - IPv6

  - MPLS

- CFP Multi-Source Agreement (MSA)-compliant management data input/output (MDIO) control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) (in all PIC and chassis configurations).
- Interface creation:
  - When the PIC is brought online, the router creates one interface, et-*x*/*y*/0, where *x* represents the FPC slot number and *y* represents PIC slot number. The physical interface represents internal Ethernet Packet Forwarding Engines.
  - The FPC slot number ranges from 0 through 7 in T640, T1600, and T4000 routers. The PIC slot numbers are 0 and 1.
  - Packet Forwarding Engine 0 is the physical interface 0, and Packet Forwarding Engine 1 is the physical interface 1.
- 802.3 link aggregation:
  - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
  - An aggregate bundle that consists purely of 40-Gigabit Ethernet PICs supports a maximum of 40-Gigabit Ethernet links depending on the system implementation.

For Junos OS configuration information about this PIC, see "Configuring 40-Gigabit Ethernet PICs" on page 471. For hardware compatibility information, see the *T1600 PICs Supported* topic in the *T1600 Core Router* hardware guide and the *T640 PICs Supported* topic in the *T640 Core Router* hardware guide, and the *T4000 PICs Supported* topic in the *T4000 Core Router* hardware guide.

Related Documentation

- Configuring 40-Gigabit Ethernet PICs on page 471
- *T640 Core Router*
- *T1600 Core Router*
- *T4000 Core Router*
- *TX Matrix Plus Router*
- *T640 PICs Supported*
- *T1600 PICs Supported*
- *T4000 PICs Supported*

 Copyright © 2013, Juniper Networks, Inc.

Configuring 40-Gigabit Ethernet PICs

You can configure the following features on the 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP):

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting for receive (**Rx**) and transmit (**Tx**)
- Multiple tag protocol ID (TPID) support
- Channels defined by two stacked VLAN tags
- Channels defined by **flex-vlan-tagging**
- IP service for stacked VLAN tags
- IP service for nonstandard TPID

The following features are not supported on the 40-Gigabit Ethernet PIC with CFP:

- MAC learning
- MAC policing
- Layer 2 rewrite

NOTE: Each 40-Gigabit Ethernet PIC with CFP creates a single et- physical interface in the Routing Engine and Packet Forwarding Engine.

The 40-Gigabit Ethernet PIC with CFP supports aggregated Ethernet configuration to achieve higher throughput capability, whereby the configuration is similar to the 1-Gigabit or 10-Gigabit aggregated Ethernet interface configuration. A maximum of 40-Gigabit Ethernet PIC links can be bundled into a single aggregated Ethernet configuration depending on the system implementation.

To configure the 40-Gigabit Ethernet PIC with CFP:

1. Perform the media configuration.

   The command used to configure the media for the 40-Gigabit Ethernet PIC with CFP is the same as that for other Ethernet PICs, such as the 4-port 10-Gigabit Ethernet PIC.

2. Specify the logical interfaces.

A single physical interface is created when the 40-Gigabit Ethernet PIC with CFP is brought online (et-*x*/*y*/0, where *x* represents the FPC slot number and *y* represents the PIC slot number). For more information, see "Configuring a Logical Interface for Access Mode" on page 75 and "Configuring a Logical Interface for Trunk Mode" on page 76.

3. Configure the 802.3 link aggregation.

   - You must explicitly configure an aggregated interface on the 40-Gigabit Ethernet PIC with CFP that includes the 40-Gigabit Ethernet interfaces. For more information, see "Configuring an Aggregated Ethernet Interface" on page 87.
   - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with the configuration of the 1-Gigabit Ethernet PIC, 10-Gigabit Ethernet PIC, and the aggregated Ethernet interfaces. In each aggregated bundle, Junos OS supports a maximum of 40-Gigabit Ethernet links. For more information, see "Configuring an Aggregated Ethernet Interface" on page 87 and "10-port 10-Gigabit Ethernet LAN/WAN PIC Overview" on page 427.

4. Configure the Packet Forwarding Engine features.

   The 40-Gigabit Ethernet PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see "Configuring Gigabit Ethernet Policers" on page 397, "Configuring MAC Address Filtering" on page 401, and "Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview" on page 197.

Related Documentation

- 40-Gigabit Ethernet PIC Overview on page 469
- Configuring Gigabit Ethernet Policers on page 397
- Configuring MAC Address Filtering on page 401
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 197

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 29

Configuring Ethernet Interfaces for PTX Series Packet Transport Routers

Understanding Ethernet Interfaces for PTX Series Packet Transport Routers

PTX Series Packet Transport Routers are a portfolio of high-performance platforms designed for the service provider supercore. A PTX Series Packet Transport Router working in conjunction with a T Series core router allows a service provider to build a core network that is flexible enough to accommodate cloud-delivered services, mobility for devices and users, and bandwidth-intensive applications such as HD video. Forwarding architecture for PTX Series Packet Transport Routers is focused on MPLS and Ethernet.

Figure 31: PTX5000 in a Juniper Networks Environment

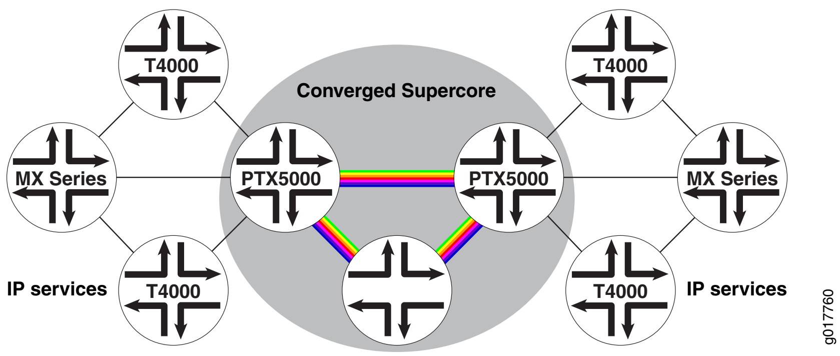

All physical interfaces on a PTX Series Packet Transport Router use *et* for the FPC type. For information about how to specify interfaces, see these topics:

- *Physical Part of an Interface Name*
- *Logical Part of an Interface Name*

NOTE: Physical interfaces on PTX Series Packet Transport Routers do not support:

- VLAN rewrite for CCC encapsulation
- Source MAC learning for accounting
- MAC policing

NOTE: Wide Area Network Physical Layer Device (WAN PHY) mode is supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number P1-PTX-24-10G-W-SFPP). No other PIC supports WAN PHY on the PTX router.

Related Documentation

- *PTX Series Packet Transport Router PIC Guide*.
- *PTX Series Packet Transport Router Management Ethernet Interfaces*
- Configuring MAC Filtering on PTX Series Packet Transport Routers on page 475

 Copyright © 2013, Juniper Networks, Inc.

- Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers on page 481

Configuring MAC Filtering on PTX Series Packet Transport Routers

This topic describes how to configure MAC filtering on PTX Series Packet Transport Routers. MAC filtering enables you to specify the MAC addresses from which the Ethernet interface can receive packets.

MAC filtering support on PTX Series Packet Transport Routers includes:

- MAC source and destination address filtering for each port.
- MAC source address filtering for each physical interface.
- MAC source address filtering for each logical interface.

When you filter logical and physical interfaces, you can specify up to 1000 MAC source addresses per port.

To configure MAC source address filtering for a physical interface, include the **source-filtering** and **source-address-filter** statements at the **[edit interfaces et-*fpc/pic/port* gigether-options]** hierarchy level:

```
[edit interfaces]
et-x/y/z {
  gigether-options {
    source-filtering;
    source-address-filter {
      mac-address;
    }
  }
}
```

The **source-address-filter** statement configures which MAC source addresses are filtered. The specified physical interface drops all packets from the MAC source addresses you specify. You can specify the MAC address as ***nn:nn:nn:nn:nn:nn*** where ***n*** is a decimal digit. To specify more than one address, include multiple ***mac-address*** options in the **source-address-filter** statement.

To configure MAC source address filtering for a logical interface, include the **accept-source-mac** statement at the **[edit interfaces et-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces]
et-x/y/z {
  gigether-options {
    source-filtering;
  }
  unit logical-unit-number {
    accept-source-mac {
      mac-address mac-address;
    }
```

Copyright © 2013, Juniper Networks, Inc.

```
    }
}
```

The **accept-source-mac** statement configures which MAC source addresses are accepted on the logical interface. You can specify the MAC address as ***nn:nn:nn:nn:nn:nn*** where ***n*** is a decimal digit. To specify more than one address, include multiple **mac-address *mac-address*** options in the **accept-source-mac** statement.

After an interface filter is configured, there is an accounting entry that is associated with the MAC address filter. Counters accumulate if there are packets with matching MAC source addresses. You can use the **show interfaces mac-database** Junos OS CLI command to view the address count.

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474
- *show interfaces mac-database (Gigabit Ethernet)*

Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers

This topic describes how to configure flexible VLAN tagging on PTX Series Packet Transport Routers. In addition to VLAN tagging and stacked VLAN tagging, you can configure a port for flexible tagging. With flexible VLAN tagging, you can configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

To configure mixed tagging, include the **flexible-vlan-tagging** statement at the **[edit interfaces et-*fpc/pic/port*]** hierarchy level. You must also include the **vlan-tags** statement with **inner** and **outer** options or the **vlan-id** statement at the **[edit interfaces et-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces et-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
   vlan-id number;
}
unit logical-unit-number {
   vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
}
```

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474

Configuring Tag Protocol IDs (TPIDs) on PTX Series Packet Transport Routers

This topic describes how to configure the TPIDs expected to be sent or received on a particular VLAN for PTX Series Packet Transport Routers.

For other types of Juniper Networks Ethernet PICs, you could configure 8 TPIDs per port. However, the PTX Series Packet Transport Routers use MTIP and TL to classify a specific TPID and Ethernet type. For MTIP, you can configure a maximum of 8 TPIDs for each MAC chip.

 Copyright © 2013, Juniper Networks, Inc.

As a consequence, you can specify the **tag-protocol-id** configuration statement only for the first port (**0**) of a PTX Series Ethernet PIC. If you configure **tag-protocol-id** statements on the other port, the configuration is ignored and a system error is recorded.

For example, the following is a supported configuration:

```
[edit interfaces et-2/0/0]
gigether-options {
  ethernet-switch-profile {
    tag-protocol-id [0x8100 0x9100];
  }
}
```

The **tag-protocol-id** configuration statement supports up to eight TPIDs on port **0** of a given Ethernet PIC. All eight TPIDs are populated to the two MTIPs and TLs associated with the Ethernet PIC.

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 476

Configuring Interface Encapsulation on PTX Series Packet Transport Routers

This topic describes how to configure interface encapsulation on PTX Series Packet Transport Routers. Use the **flexible-ethernet-services** configuration statement to configure different encapsulation for different logical interfaces under a physical interface. With flexible Ethernet services encapsulation, you can configure each logical interface encapsulation without range restrictions for VLAN IDs.

Supported encapsulations for physical interfaces include:

- **flexible-ethernet-services**
- **ethernet-ccc**
- **ethernet-tcc**

Supported encapsulations for logical interfaces include:

- **ENET2**
- **vlan-ccc**
- **vlan-tcc**

NOTE: PTX Series Packet Transport Routers do not support **extended-vlan-cc** and **extended-vlan-tcc** encapsulation on logical interfaces. Instead, you can configure a tag protocol ID (TPID) value of 0x9100 to achieve the same results.

To configure flexible Ethernet services encapsulation, include the **encapsulation flexible-ethernet-services** statement at the **[edit interfaces et-*fpc/pic/port*]** hierarchy level. For example:

```
interfaces {
   et-fpc/pic/port {
      vlan-tagging;
      encapsulation flexible-ethernet-services;
      unit 0 {
         vlan-id 1000;
         family inet {
            address 11.0.0.20/24;
         }
      }
      unit 1 {
         encapsulation vlan-ccc;
         vlan-id 1010;
      }
      unit 2 {
         encapsulation vlan-tcc;
         vlan-id 1020;
         family tcc {
            proxy {
               inet-address 11.0.2.160;
            }
            remote {
               inet-address 11.0.2.10;
            }
         }
      }
   }
}
```

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474

Configuring Ethernet 802.3ah OAM on PTX Series Packet Transport Routers

The IEEE 802.3ah standard for Operation, Administration, and Management (OAM) provides a specification for *Ethernet in the first mile (EFM)* connectivity. EFM defines how Ethernet can be transmitted over new media types using new Ethernet physical layer (PHY) interfaces. You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah OAM standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology.

For Ethernet interfaces capable of running at 100 Mbps or faster, the IEEE 802.3ah OAM standard is supported on numerous Juniper Networks routers and switches. This topic describes configuration support for IEEE 802.3ah OAM features on PTX Series Packet Transport Routers.

 Copyright © 2013, Juniper Networks, Inc.

On PTX Series Packet Transport Routers, Junos OS Release 12.1 supports the following IEEE 802.3ah OAM features at the physical interface level:

- Discovery and link monitoring
- Fault signaling and detection
- Periodic packet management (PPM) processing
- Action profile support
- Graceful Routing Engine switchover (GRES)

To configure 802.3ah OAM support for Ethernet interfaces, include the **oam** statement at the **[edit protocols]** hierarchy level:

```
oam {
  ethernet {
    link-fault-management {
      interfaces {
        interface-name {
          pdu-interval interval;
          link-discovery (active | passive);
          pdu-threshold count;
        }
      }
    }
  }
}
```

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 373
- Configuring Link Discovery on page 374
- Detecting Remote Faults on page 379
- Configuring an OAM Action Profile on page 380

Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers

The IEEE 802.1ag provides a specification for Ethernet connectivity fault management (CFM). The Ethernet network may be comprised of one or more service instances. A service instance could be a VLAN, or a concatenation of VLANs. The goal of CFM is to provide a mechanism to monitor, locate, and isolate faulty links. Ethernet 802.1ag is supported on numerous Juniper Networks routers and switches. This topic describes configuration support for Ethernet OAM 802.1ag features on the PTX Series Packet Transport Routers.

Copyright © 2013, Juniper Networks, Inc.

Supported features include:

- Maintenance domain (**maintenance-domain** ***domain-name***) and maintenance levels (**level** ***number***).
- Maintenance association (**maintenance-association** ***ma-name***), including name formats (**name-format** and **short-name-format** for **vlan** and **2octet**), loss threshold (**loss-threshold** ***number***), and hold interval (**hold-interval** ***minutes***).
- maintenance association endpoint (MEP) functions, including Maintenance Endpoint ID (**mep** ***mep-id***), direction down (**direction down**), and autodiscovery (**auto-discovery**).
- Link trace for down MEPs (**link-down**).
- action profile (**action-profile** ***profile-name***)
- Loopback message generation and reply for down MEPs.

Features that are not supported include:

- Up MEP configuration.
- maintenance association intermediate point (MIP) configuration.

To configure flexible Ethernet services encapsulation on PTX Series Packet Transport Routers, include the **oam** statement at the **[edit protocols]** hierarchy level. For example:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md1 {
        level 0;
        maintenance-association ma1 {
          continuity-check {
            interval 100ms;
          }
          mep 1 {
            interval et-0/1/1;
            direction down;
            auto-discovery;
          }
        }
      }
    }
  }
}
```

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474
- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 249

Copyright © 2013, Juniper Networks, Inc.

Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a link aggregation group (LAG) or bundle. Link aggregation can be used for point-to-point connections. It balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

This topic describes how to configure aggregated Ethernet interfaces on PTX Series Packet Transport Routers.

On PTX Series Packet Transport Routers, aggregated Ethernet support includes the following features:

- A consistent interface type (**et** ***fpc/pic/port***) across all Ethernet interfaces.
- Ability to bundle multiple Ethernet interfaces
- Fault tolerance
- Load balancing between child links
- Advanced features including flexible VLAN tagging and Ethernet services encapsulation

Aggregated Ethernet interfaces can use interfaces from different FPCs or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
  aggregated-devices {
  ethernet {
    device-count 2;
  }
}
```

```
[edit interfaces]
et-0/0/0 {
  gigether-options {
    802.3ad ae0;
    }
  }
  et-0/0/1 {
    gigether-options {
      802.3ad ae0;
    }
  }
  ae0 {
    vlan-tagging;
    unit 0 {
      vlan-id 100;
      family inet {
        address 200.200.1.2/24;
      }
```

```
    }
    unit 1 {
        vlan-id 101;
        family inet {
            address 200.200.2.2/24;
        }
    }
}
```

Related Documentation

- Understanding Ethernet Interfaces for PTX Series Packet Transport Routers on page 474
- Configuring Junos OS for Supporting Aggregated Devices on page 88

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 30

Configuring Point-to-Point Protocol over Ethernet

Copyright © 2013, Juniper Networks, Inc.

PPPoE Overview

The Point-to-Point Protocol over Ethernet (PPPoE) connects multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

A J Series router can be configured as the CPE device for PPPoE connections. To use PPPoE, you must configure the router as a PPPoE client, encapsulate PPP packets over Ethernet, and initiate a PPPoE session.

NOTE: **J4300 and J6300 routers with asymmetrical DSL (ADSL) Physical Interface Modules (PIMs) and symmetrical high-speed DSL (SHDSL) PIMs can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections. For information about configuring ADSL and SHDSL interfaces, see *ATM-over-ADSL Overview* and *ATM-over-SHDSL Overview*.**

M120, M320, and MX Series routers can be configured as a PPPoE access concentrator server. To configure a PPPoE server on an M120, M320, or MX Series Ethernet logical interface, specify PPPoE encapsulation, include the **pp0** statement for the pseudo PPPoE physical interface, and include the **server** statement in the PPPoE options under the logical interface.

NOTE: **PPPoE encapsulation is not supported on M120, M320, or MX Series routers on an ATM2 IQ interface.**

On the J Series router, PPPoE establishes a point-to-point connection between the client (the Services Router) and the server, also called an access concentrator. Multiple hosts can be connected to the Services Router, and their data can be authenticated, encrypted, and compressed before the traffic is sent to the PPPoE session on the Services Router's Fast Ethernet or ATM-over-ADSL interface. PPPoE is easy to configure and enables services to be managed on a per-user basis rather than on a per-site basis.

This overview contains the following topics:

PPPoE Interfaces

The PPPoE interface to the access concentrator can be a Fast Ethernet interface on any Services Router, a Gigabit Ethernet interface on J4350 and J6350 Services Routers, an ATM-over-ADSL or ATM-over-SHDSL interface on all J Series Services Routers except the J2300, or an ATM-over-SHDSL interface on a J2300 Services Router. The PPPoE

 Copyright © 2013, Juniper Networks, Inc.

configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

- If the interface is Fast Ethernet, use a PPPoE encapsulation.
- If the interface is ATM over ADSL, use a PPPoE over ATM encapsulation.

The PPPoE interface on M120 or M320 routers acting as a access concentrator can be a Gigabit Ethernet or 10-Gigabit Ethernet interface.

Ethernet Interface

The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 32 on page 485 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 32: PPPoE Session on an Ethernet Loop

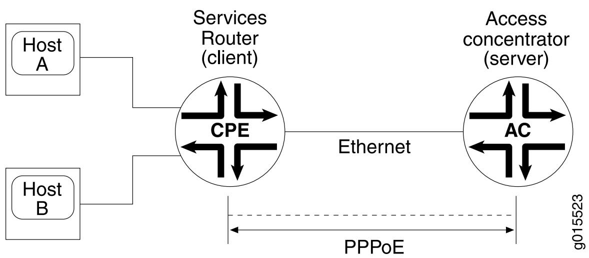

PPPoE Stages

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the discovery stage, the client discovers the access concentrator by identifying the Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the PPPoE session stage, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

NOTE: If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.

If you do not configure the access concentrator name, the PPPoE session starts using any available server in the network.

PPPoE Discovery Stage

A Services Router initiates the PPPoE discovery stage by broadcasting a PPPoE active discovery initiation (PADI) packet. To provide a point-to-point connection over Ethernet, each PPPoE session must learn the Ethernet MAC address of the access concentrator and establish a session with a unique session ID. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.

Copyright © 2013, Juniper Networks, Inc.

NOTE: A Services Router cannot receive PPPoE packets from two different access concentrators on the same physical interface.

The PPPoE discovery stage consists of the following steps:

1. PPPoE active discovery initiation (PADI)—The client initiates a session by broadcasting a PADI packet on the LAN to request a service.
2. PPPoE active discovery offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains it own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.
3. PPPoE active discovery request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.
4. PPPoE active discovery session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session.
   - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.
   - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

PPPoE Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends the PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A Services Router supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions on all interfaces on the Services Router.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE active discovery termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic.

Optional CHAP Authentication

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

Copyright © 2013, Juniper Networks, Inc.

If you configure an interface to handle incoming CHAP packets only (by including the **passive** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level), the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not include the **passive** statement, the interface always challenges its peer.

For more information about CHAP, see *Configuring the PPP Challenge Handshake Authentication Protocol*.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Developing a Log Storage Strategy*
- Evaluation Order for Matching Client Information in PPPoE Service Name Tables on page 493
- Benefits of Configuring PPPoE Service Name Tables on page 493
- Configuring PPPoE on page 494
- Disabling the Sending of PPPoE Keepalive Messages on page 502
- Configuring PPPoE Service Name Tables on page 502
- Creating a Service Name Table on page 503
- Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag on page 504
- Configuring the Action Taken for the Any Service on page 505
- Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag on page 506
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 507
- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 508
- Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client on page 509
- Enabling Advertisement of Named Services in PADO Control Packets on page 510
- Assigning a Service Name Table to a PPPoE Underlying Interface on page 510
- Example: Configuring a PPPoE Service Name Table on page 512
- Tracing PPPoE Operations on page 514
- Troubleshooting PPPoE Service Name Tables on page 516
- Verifying a PPPoE Configuration on page 518
- *Ethernet Interfaces*

Understanding PPPoE Service Name Tables

On an M120 router, M320 router, or MX Series router acting as a remote access concentrator (AC), also referred to as a *PPPoE server*, you can configure up to 32 PPPoE service name tables and assign the service name tables to PPPoE underlying interfaces. A *PPPoE service name table* defines the set of *services* that the router can provide to a PPPoE client. Service entries configured in a PPPoE service name table represent the *service name tags* transmitted between the client and the router in a PPPoE control packet.

This overview covers the following topics to help you understand and configure PPPoE service name tables:

Interaction Among PPPoE Clients and Routers During the Discovery Stage

In networks with mesh topologies, PPPoE clients are often connected to multiple PPPoE servers (remote ACs). During the PPPoE discovery stage, a PPPoE client identifies the Ethernet MAC address of the remote AC that can service its request, and establishes a unique PPPoE session identifier for a connection to that AC.

The following steps describe, at a high level, how the PPPoE client and the remote AC (router) use the PPPoE service name table to interact during the PPPoE discovery stage:

1. The PPPoE client broadcasts a PPPoE Active Discovery Initiation (PADI) control packet to all remote ACs in the network to request that an AC support certain services.

   The PADI packet must contain either, but not both, of the following:

   - One and only one nonzero-length service name tag that represents a specific client service
   - One and only one empty (zero-length) service name tag that represents an unspecified service

2. One or more remote ACs respond to the PADI packet by sending a PPPoE Active Discovery Offer (PADO) packet to the client, indicating that the AC can service the client request.

   To determine whether it can service a particular client request, the router matches the service name tag received in the PADI packet against the service name tags configured in its service name table. If a matching service name tag is found in the PPPoE service name table, the router sends the client a PADO packet that includes

 Copyright © 2013, Juniper Networks, Inc.

the name of the AC from which it was sent. If no matching service name tag is found in the PPPoE service name table, the router drops the PADI request and does not send a PADO response to the client.

3. The PPPoE client sends a unicast PPPoE Active Discovery Request (PADR) packet to the AC to which it wants to connect, based on the responses received in the PADO packets.
4. The selected AC sends a PPPoE Active Discovery Session (PADS) packet to establish the PPPoE connection with the client.

Service Entries and Actions in PPPoE Service Name Tables

A PPPoE service name table can include three types of service entries: named services, an **empty** service, and an **any** service. For each service entry, you specify the action to be taken by the underlying interface when the router receives a PADI packet containing the specified service name tag.

You can configure the following services and actions in a PPPoE service name table:

- Named service—Specifies a PPPoE client service that an AC can support. For example, you might configure named services associated with different subscribers who log in to the PPPoE server, such as **user1-service** or **user2-service**, or that correspond to different ISP service level agreements, such as **premium** and **standard**. Each PPPoE service name table can include a maximum of 512 named service entries, excluding **empty** and **any** service entries. A named service is associated with the **terminate** action by default.
- **empty** service—A service tag of zero length that represents an unspecified service. Each PPPoE service name table includes one empty service. The **empty** service is associated with the **terminate** action by default.
- **any** service—Acts as a default service for non-empty service entries that do not match the named service entries or **empty** service entry configured in the PPPoE service name table. Each PPPoE service name table includes one **any** service. The **any** service is useful when you want to match the agent circuit identifier and agent remote identifier information for a PPPoE client, but do not care about the contents of the service name tag transmitted in the control packet. The **any** service is associated with the **drop** action by default.
- Action—Specifies the action taken by the underlying PPPoE interface assigned to the PPPoE service name table on receipt of a PADI packet from the client containing a particular service request. You can configure one of the following actions for the associated named service, **empty** service, **any** service, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in the PPPoE service name table on the router:
    - **terminate**—(Default) Directs the router to immediately respond to the PADI packet by sending the client a PADO packet containing the name of the AC that can service the request. Named services, **empty** services, and ACI/ARI pairs are associated with the **terminate** action by default. Configuring the **terminate** action for a service enables you to more tightly control which PPPoE clients can access and receive services from a particular PPPoE server.

- **delay**—Number of seconds that the PPPoE underlying interface waits after receiving a PADI packet from the client before sending a PADO packet in response. In networks with mesh topologies, you might want to designate a primary PPPoE server and a backup PPPoE server for handling a particular service request. In such a scenario, you can configure a delay for the associated service entry on the backup PPPoE server to allow sufficient time for the primary PPPoE server to respond to the client with a PADO packet. If the primary server does not send the PADO packet within the delay period configured on the backup server, then the backup server sends the PADO packet after the delay period expires.
- **drop**—Directs the router to drop (ignore) a PADI packet containing the specified service name tag when received from a PPPoE client, which effectively denies the client's request to provide the associated service. The **any** service is associated with the **drop** action by default. To prohibit the router from responding to PADI packets that contain **empty** or **any** service name tags, you can configure the **drop** action for the empty or **any** service. You can also use the **drop** action in combination with ACI/ARI pairs to accept specific service name tags only from specific subscribers, as described in the following information about ACI/ARI pairs.

ACI/ARI Pairs in PPPoE Service Name Tables

To specify agent circuit identifier (ACI) and agent remote identifier (ARI) information for a named service, **empty** service, or **any** service in a PPPoE service name table, you can configure an ACI/ARI pair. An ACI/ARI pair contains an agent circuit ID string that identifies the DSLAM interface that initiated the service request, and an agent remote ID string that identifies the subscriber on the DSLAM interface that initiated the service request. You can think of an ACI/ARI pair as the representation of one or more PPPoE clients accessing the router by means of the PPPoE service name table.

ACI/ARI specifications support the use of wildcard characters in certain formats. You can configure a combined maximum of 8000 ACI/ARI pairs, both with and without wildcards, per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the service entries in the service name table.

You must specify the action—**terminate**, **delay**, or **drop**—taken by the underlying PPPoE interface when it receives a client request containing vendor-specific ACI/ARI information that matches the ACI/ARI information configured in the PPPoE service name table on the router. An ACI/ARI pair is associated with the **terminate** action by default.

For example, assume that for the **user1-service** named service, you configure the **drop** action for the service and the **terminate** action for the associated ACI/ARI pairs. In this case, the ACI/ARI pairs identify the DSLAM interfaces and associated subscribers authorized to access the PPPoE server. Using this configuration causes the router to drop PADI packets containing the **user1-service** tag *unless* the PADI packet also contains vendor-specific ACI/ARI information that matches the subscribers identified in one or more of the ACI/ARI pairs. For PADI packets containing matching ACI/ARI information, the router sends an immediate PADO response to the client indicating that it can provide the requested service for the specified subscribers.

 Copyright © 2013, Juniper Networks, Inc.

You can also associate a PPPoE dynamic profile, routing instance, and static PPPoE interface with an ACI/ARI pair.

Dynamic Profiles and Routing Instances in PPPoE Service Name Tables

You can associate a previously configured PPPoE dynamic profile with a named service, **empty** service, or **any** service in the PPPoE service name table, or with an ACI/ARI pair defined for these services. The router uses the attributes defined in the profile to instantiate a dynamic PPPoE subscriber interface based on the service name, ACI, and ARI information provided by the PPPoE client during PPPoE negotiation. The dynamic profile configured for a service entry or ACI/ARI pair in a PPPoE service name table overrides the dynamic profile assigned to the PPPoE underlying interface on which the dynamic PPPoE interface is created.

To specify the routing instance in which to instantiate the dynamic PPPoE interface, you can associate a previously configured routing instance with a named service, **empty** service, or **any** service in the PPPoE service name table, or with an ACI/ARI pair defined for these services. Like dynamic profiles configured for service entries or ACI/ARI pairs, the routing instance configured for the PPPoE service name table overrides the routing instance assigned to the PPPoE underlying interface.

For information about configuring the PPPoE service name table to create a dynamic PPPoE subscriber interface, see *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation* in the *Junos OS Subscriber Management and Services Library*.

Maximum Sessions Limit in PPPoE Service Name Tables

To limit the number of PPPoE client sessions that can use a particular service entry in the PPPoE service name table, you can configure the maximum number of active PPPoE sessions using either dynamically-created or statically-created PPPoE interfaces that the router can establish with a particular named service, **empty** service, or **any** service. (You cannot configure the maximum sessions limit for an ACI/ARI pair.) The maximum sessions limit must be in the range 1 through the platform-specific maximum PPPoE sessions supported for your routing platform. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

The router uses the maximum sessions value for a service entry in the PPPoE service name table in conjunction with both of the following:

- The maximum sessions (**max-sessions**) value configured for the PPPoE underlying interface
- The maximum number of PPPoE sessions supported on your routing platform

If your configuration exceeds either of these maximum session limits, the router cannot establish the PPPoE session.

Copyright © 2013, Juniper Networks, Inc.

Static PPPoE Interfaces in PPPoE Service Name Tables

To reserve a previously configured static PPPoE interface for use only by the PPPoE client with matching ACI/ARI information, you can specify a single static PPPoE interface for each ACI/ARI pair defined for a named service entry, **empty** service entry, or **any** service entry in a PPPoE service name table. (You cannot configure a static interface for a service entry that does not have an ACI/ARI pair defined.) The static PPPoE interface associated with an ACI/ARI pair takes precedence over the general pool of static PPPoE interfaces associated with the PPPoE underlying interface configured on the router.

When you configure a static interface in the PPPoE service name table, make sure there is a one-to-one correspondence between the PPPoE client and the static interface. For example, if two clients have identical ACI/ARI information that matches the information in the PPPoE service name table, the router reserves the static interface for exclusive use by the first client that logs in to the router. As a result, the router prevents the second client from logging in.

NOTE: **You cannot configure a static interface for an ACI/ARI pair already configured with a dynamic profile and routing instance. Conversely, you cannot configure a dynamic profile and routing instance for an ACI/ARI pair already configured with a static interface.**

PADO Advertisement of Named Services in PPPoE Service Name Tables

By default, the advertisement of named services in PADO control packets sent by the router to the PPPoE client is disabled. You can enable advertisement of named services in the PADO packet as a global option when you configure the PPPoE protocol on the router. Configuring PADO advertisement notifies PPPoE clients of the services that the router (AC) can offer.

If you enable advertisement of named services in PADO packets, make sure the number and length of all advertised service entries does not exceed the maximum transmission unit (MTU) size supported by the PPPoE underlying interface.

Related Documentation

- Evaluation Order for Matching Client Information in PPPoE Service Name Tables on page 493
- Benefits of Configuring PPPoE Service Name Tables on page 493
- Configuring PPPoE Service Name Tables on page 502
- Example: Configuring a PPPoE Service Name Table on page 512

- For information about creating dynamic PPPoE subscriber interfaces, see *Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles* in the *Junos OS Subscriber Management and Services Library*

- PPPoE Overview on page 484

- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Evaluation Order for Matching Client Information in PPPoE Service Name Tables

When the router receives a service request from a PPPoE client, it evaluates the entries configured in the PPPoE service name table to find a match for the client's ACI/ARI information so it can take the appropriate action.

The order of evaluation is as follows:

1. The router evaluates the ACI/ARI information configured for the **any** service entry, and ignores the contents of the service name tag transmitted by the client.
2. If no match is found for the client information, the router evaluates the ACI/ARI information for the **empty** service entry and the named service entries. If an ACI/ARI pair is not configured for these service entries, the router evaluates the other attributes configured for the **empty** service and named services.
3. If there is still no match for the client information, the router evaluates the other attributes configured for the **any** service entry, and ignores both the ACI/ARI information for the **any** service and the contents of the service name tag transmitted by the client. If the **any** service is configured for the default action, **drop**, the router drops the PADR packet. If the **any** service is configured for a nondefault action (**terminate** or **delay**), the router evaluates the other attributes configured for the **any** service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Benefits of Configuring PPPoE Service Name Tables on page 493
- Configuring PPPoE Service Name Tables on page 502
- *Example: Configuring a PPPoE Service Name Table for Dynamic Subscriber Interface Creation*
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Benefits of Configuring PPPoE Service Name Tables

This topic describes the benefits of configuring PPPoE service name tables.

Configuring PPPoE service name tables provides the following benefits:

- Enables support for multiple services requested by PPPoE clients, and configuration of an action for the underlying PPPoE interface to take (**delay**, **drop**, or **terminate**) upon receipt of a PPPoE Active Discovery Initiation (PADI) packet requesting that service.
- Provides tighter control over which PPPoE clients can log in to and receive services from a particular PPPoE server.
- Provides load balancing across a set of remote access concentrators (ACs) in a mesh topology by enabling you to configure agent circuit identifier/agent remote identifier (ACI/ARI) pairs for named, **empty**, and **any** service entries to specify the appropriate AC to receive and service a particular PPPoE client request.

Copyright © 2013, Juniper Networks, Inc.

- Offers a more targeted approach to configuration of PPPoE sessions based on the service name and ACI/ARI information provided by the PPPoE client during PPPoE negotiation.
- Supports creation of dynamic PPPoE subscriber interfaces in a specified routing instance based on configuration of a service entry or ACI/ARI pair in the PPPoE service name table.
- Enables you to reserve a specified static PPPoE interface for use only by the PPPoE client with matching ACI/ARI information.
- Enables you to specify the maximum number of PPPoE client sessions that can use a particular service entry in the PPPoE service name table.
- Provides redundancy across a set of remote ACs in a mesh topology by enabling you to configure a primary AC and a backup AC for handling a specific service request from a PPPoE client.

  For example, on the primary AC for handling a client service, you might configure the **terminate** action for the associated service to direct the primary AC to immediately send a PPPoE Active Discovery Offer (PADO) packet in response to a PADI packet containing that service name tag. On the backup AC for the client service, you might configure the **delay** action for the associated service to specify the number of seconds the backup AC waits after receiving a PADI packet from the client before sending a PADO packet in response. If the primary AC does not send a PADO packet to the client within the delay period configured on the backup AC, then the backup AC sends the PADO packet after the delay period expires.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- Example: Configuring a PPPoE Service Name Table on page 512
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Configuring PPPoE

To configure PPPoE on a J Series Services Router, perform the following tasks:

1. Configure PPPoE encapsulation for an Ethernet interface or Ethernet over ATM encapsulation for an ATM-over-ADSL interface.
2. If you are configuring ATM over ADSL, configure LLC encapsulation on the logical interface.
3. Specify the logical Ethernet interface or the logical ATM interface as the underlying interface for the PPPoE session.
4. Configure the operational mode as client.
5. Identify the access concentrator by a unique name.

 Copyright © 2013, Juniper Networks, Inc.

6. Optionally, specify how many seconds to wait before attempting to reconnect.
7. Provide a name for the type of service provided by the access concentrator.
8. Optionally, configure the maximum transmission unit (MTU) of the interface.
9. Configure the PPPoE interface address.
10. Configure the destination PPPoE interface address.
11. Optionally, configure the MTU size for the protocol family.
12. Optionally, disable the sending of keepalive messages on the logical interface.

To configure PPPoE on an M120 or M320 Multiservice Edge Router or MX Series Universal Edge Router operating as an access concentrator, perform the following tasks:

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
3. Optionally, configure the maximum transmission unit (MTU) of the interface.
4. Configure the operational mode as server.
5. Configure the PPPoE interface address.
6. Configure the destination PPPoE interface address.
7. Optionally, configure the MTU size for the protocol family.
8. Optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.
9. Optionally, disable the sending of PADS messages that contain certain error tags.

Setting the Appropriate Encapsulation on the PPPoE Interface

For PPPoE on an Ethernet interface, you must configure encapsulation on the logical interface and use PPP over Ethernet encapsulation.

For PPPoE on an ATM-over-ADSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL logical interface, use PPPoE over AAL5 LLC encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.

NOTE: **PPPoE encapsulation is not supported on an M120 or M320 router on an ATM2 IQ interface.**

When you configure a point-to-point encapsulation such as PPP on a physical interface, the physical interface can have only one logical interface (only one **unit** statement) associated with it.

To configure physical interface properties, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces interface-name]
encapsulation ethernet-over-atm;
```

To configure logical interface encapsulation properties, include the **encapsulation** statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Perform the task appropriate for the interface on which you are using PPPoE:

- Configuring PPPoE Encapsulation on an Ethernet Interface on page 496
- Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 496

Configuring PPPoE Encapsulation on an Ethernet Interface

Both the client and the server must be configured to support PPPoE. To configure PPPoE encapsulation on an Ethernet interface, include the **encapsulation** statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]**

Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

To configure the PPPoE encapsulation on a ATM-over-ADSL interface, perform the following steps:

1. Include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level, and specify **ethernet-over-atm**:

   ```
   [edit interfaces pp0]
   encapsulation ethernet-over-atm;
   ```

2. Configure LLC encapsulation on the logical interface by including the **encapsulation** statement and specifying **ppp-over-ether-over-atm-llc**:

   ```
   encapsulation ppp-over-ether-over-atm-llc;
   ```

You can include this statement at the following hierarchy levels:

 Copyright © 2013, Juniper Networks, Inc.

- [edit interfaces pp0 unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]

Configuring a PPPoE Interface

NOTE: When you configure a static PPPoE logical interface, you must include the **pppoe-options** subhierarchy at the **[edit interfaces pp0 unit *logical-unit-number*]** hierarchy level or at the **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]** hierarchy level. If you omit the **pppoe-options** subhierarchy from the configuration, the commit operation fails.

Configuring the PPPoE Underlying Interface

To configure the underlying Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface, include the **underlying-interface** statement:

```
underlying-interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Specify the logical Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface as the underlying interface—for example, **at-0/0/1.0** (ATM VC), **fe-1/0/1.0** (Fast Ethernet interface), or **ge-2/0/0** (Gigabit Ethernet interface).

Copyright © 2013, Juniper Networks, Inc.

Identifying the Access Concentrator

When configuring a PPPoE client, identify the access concentrator by a unique name by including the **access-concentrator** statement:

access-concentrator *name*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Automatic Reconnect Wait Timer

By default, after a PPPoE session is terminated, the session attempts to reconnect immediately. When configuring a PPPoE client, you can specify how many seconds to wait before attempting to reconnect, by including the **auto-reconnect** statement:

auto-reconnect *seconds*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

You can configure the reconnection attempt to occur in 0 through 4,294,967,295 seconds after the session terminates.

Configuring the PPPoE Service Name

When configuring a PPPoE client, identify the type of service provided by the access concentrator—such as the name of the Internet service provider (ISP), class, or quality of service—by including the **service-name** statement:

service-name *name*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Server Mode

When configuring a PPPoE server, identify the mode by including the **server** statement:

server;

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**

 Copyright © 2013, Juniper Networks, Inc.

- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Client Mode

When configuring a PPPoE client, identify the mode by including the **client** statement:

```
client;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number* pppoe-options]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]**

Configuring the PPPoE Source and Destination Addresses

When configuring a PPPoE client or server, assign source and destination addresses—for example, **192.168.1.1/32** and **192.168.1.2**. To assign the source and destination address, include the **address** and **destination** statements:

```
address address {
  destination address;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces pp0.0 family inet]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family inet]**

Deriving the PPPoE Source Address from a Specified Interface

For a router supporting PPPoE, you can derive the source address from a specified interface—for example, the loopback interface, **lo0.0**—and assign a destination address—for example, **192.168.1.2**. The specified interface must include a logical unit number and have a configured IP address. To derive the source address and assign the destination address, include the **unnumbered-address** and **destination** statements:

```
unnumbered-address interface-name destination address;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces pp0.0 family inet]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family inet]**

Configuring the PPPoE IP Address by Negotiation

You can have the PPPoE client router obtain an IP address by negotiation with the remote end. This method might require the access concentrator to use a RADIUS authentication server. To obtain an IP address from the remote end by negotiation, include the **negotiate-address** statement:

```
negotiate-address;
```

Copyright © 2013, Juniper Networks, Inc.

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0.0 family (inet | inet6 | mpls)]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family (inet | inet6 | mpls**)]

Configuring the Protocol MTU PPPoE

You can configure the maximum transmission unit (MTU) size for the protocol family. Specify a range from 0 through 5012 bytes. Ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. To set the MTU, include the **mtu** statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0.0 family (inet | inet6 | mpls)]**
- **[edit logical-systems *logical-system-name* interfaces pp0.0 family (inet | inet6 | mpls**)]

You can modify the MTU size of the interface by including the **mtu *bytes*** statement at the **[edit interfaces pp0]** hierarchy level:

```
[edit interfaces pp0]
mtu bytes;
```

The default media MTU size used and the range of available sizes on a physical interface depends on the encapsulation used on that interface.

Example: Configuring a PPPoE Client Interface on a J Series Services Router

Configure a PPPoE over ATM-over-ADSL interface:

```
[edit interfaces]
at-2/0/0 {
  encapsulation ethernet-over-atm;
  atm-options {
    vpi 0;
  }
  dsl-options {
    operating-mode auto;
  }
  unit 0 {
    encapsulation ppp-over-ether-over-atm-llc;
    vci 0.120;
  }
}
pp0 {
  mtu 1492;
  unit 0 {
    ppp-options {
      chap {
        access-profile A-ppp-client;
        local-name A-at-2/0/0.0;
      }
    }
```

 Copyright © 2013, Juniper Networks, Inc.

```
            pppoe-options {
                underlying-interface at-2/0/0;
                client;
                access-concentrator ispl.com;
                service-name "video@ispl.com";
                auto-reconnect 100;
            }
            no-keepalives;
            family inet {
                negotiate-address;
                mtu 100;
            }
            family inet6 {
                negotiate-address;
                mtu 200;
            }
            family mpls {
                negotiate-address;
                mtu 300;
            }
        }
    }
```

Example: Configuring a PPPoE Server Interface on an M120 or M320 Router

Configure a PPPoE server over a Gigabit Ethernet interface:

```
[edit interfaces]
ge-1/0/0 {
    vlan-tagging;
    unit 1 {
        encapsulation ppp-over-ether;
        vlan-id 10;
    }
}
pp0 {
    unit 0 {
        pppoe-options {
            underlying-interface ge-1/0/0.0;
            server;
        }
        ppp-options {
        }
        family inet {
            address 22.2.2.1/32 {
                destination 22.2.2.2;
            }
        }
    }
}
```

Related Documentation

- PPPoE Overview on page 484
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Disabling the Sending of PPPoE Keepalive Messages

When configuring the client, you can disable the sending of keepalive messages on a logical interface by including the **no-keepalives** statement:

```
no-keepalives;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces pp0 unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]**

Related Documentation

- PPPoE Overview on page 484
- *Ethernet Interfaces*

Configuring PPPoE Service Name Tables

To configure PPPoE service name tables:

1. Create a PPPoE service name table.

   See "Creating a Service Name Table" on page 503.
2. (Optional) Configure the action taken for the **empty** service.

   See "Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag" on page 504.
3. (Optional) Configure the action taken for the **any** service.

   See "Configuring the Action Taken for the Any Service" on page 505.
4. Assign a named service to the service name table and optionally configure the action taken for the specified service name.

   See "Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag" on page 506.
5. (Optional) Configure the action taken for an ACI/ARI pair associated with a service.

   See "Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information" on page 507.
6. (Optional) Assign a dynamic profile and routing instance to a service name or ACI/ARI pair to instantiate a dynamic PPPoE subscriber interface.

   See *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*.
7. (Optional) Limit the number of active PPPoE sessions that the router can establish with the specified service.

   See "Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name" on page 508.

Copyright © 2013, Juniper Networks, Inc.

8. (Optional) Assign a static PPPoE interface to an ACI/ARI pair to reserve the interface for exclusive use by the PPPoE client with matching ACI/ARI information.

   See "Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client" on page 509.

9. (Optional) Enable advertisement of named services in the PADO control packet sent by the router to the client.

   See "Enabling Advertisement of Named Services in PADO Control Packets" on page 510.

10. Assign a service name table to a PPPoE underlying interface.

    See "Assigning a Service Name Table to a PPPoE Underlying Interface" on page 510.

11. (Optional) Configure trace options for troubleshooting the configuration.

    See "Tracing PPPoE Operations" on page 514.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Benefits of Configuring PPPoE Service Name Tables on page 493
- Example: Configuring a PPPoE Service Name Table on page 512
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Creating a Service Name Table

You can create up to 32 PPPoE service name tables on the router. You can optionally create named services and add them to a service name table. By default, the **empty** service and the **any** service are present in each service name table.

A named service specifies a PPPoE client service that the router, functioning as an access concentrator or PPPoE server, can support. The **empty** service is a service tag of zero length that represents an unspecified service. The **any** service acts as a default service for non-empty service entries that do not match the named or **empty** service entries configured in the PPPoE service name table. Named services and the **empty** service are associated with the **terminate** action by default, and the **any** service is associated with the **drop** action by default.

To create a PPPoE service name table:

- Specify the table name.

```
[edit protocols pppoe]
user@host# set service-name-tables table1
```

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Understanding PPPoE Service Name Tables on page 488
- PPPoE Overview on page 484

- *Ethernet Interfaces*

Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag

You can configure the action taken by the PPPoE underlying interface when it receives a PADI packet that includes a zero-length (empty) service name tag. The **empty** service is present by default in every PPPoE service name table.

To indicate that it can service the client request, the interface returns a PADO packet in response to the PADI packet. By default, the interface immediately responds to the request; this is the **terminate** action. Alternatively, you can configure the **drop** action to ignore (drop) the PADI packet, or the **delay** action to set a delay between receipt of the PADI packet and transmission of the PADO packet.

(Optional) To configure the action taken for the **empty** service in response to a PADI packet from a PPPoE client:

- Specify the action.

```
[edit protocols pppoe service-name-tables table1]
user@host# set service empty drop
```

You can also accomplish the following optional tasks when you configure the **empty** service:

- Specify the agent circuit identifier (ACI) and agent remote identifier (ARI) information to determines the action taken by the PPPoE underlying interface when it receives a PADI packet with matching ACI/ARI information.
- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Limit the number of active PPPoE sessions that the router can establish with the **empty** service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 507
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*
- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 508
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Configuring the Action Taken for the Any Service

The **any** service acts as a default service for service name tags transmitted by the client that do not match any of the service entries configured in the PPPoE service name table on the router. By configuring an action for the **any** service, you specify the action taken by the PPPoE underlying interface when it receives a PADI control packet from a client that includes a non-empty service name tag that does not match any of the named service entries or **empty** service entry in the PPPoE service name table.

Each PPPoE service name table includes one **any** service entry associated by default with the **drop** action. The **drop** action ignores a PADI packet containing a nonmatching service name tag. Alternatively, you can configure the **terminate** action to immediately respond to the PADI packet with a PADO packet, or the **delay** action to specify a delay between receipt of the PADI packet and transmission of the PADO packet.

To configure the action taken for the **any** service in response to a PADI packet from a PPPoE client:

- Specify the action.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service any terminate
  ```

You can also accomplish the following optional tasks when you configure the **any** service:

- Specify the agent circuit identifier (ACI) and agent remote identifier (ARI) information to determine the action taken by the PPPoE underlying interface when it receives a PADI packet with matching ACI/ARI information.
- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Limit the number of active PPPoE sessions that the router can establish with the **any** service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 507
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*
- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 508
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag

You can configure a maximum of 512 named service entries, excluding **empty** and **any** service entries, across all PPPoE service name tables on the router. A named service specifies a PPPoE client service that the router, functioning as an access concentrator or PPPoE server, can support. You can optionally configure the action taken by the PPPoE underlying interface when it receives a PADI packet that includes a matching named service (service name tag).

To indicate that it can service the client request, the interface returns a PADO packet in response to the PADI packet. By default, the interface immediately responds to the request; this is the **terminate** action. Alternatively, you can configure the **drop** action to ignore (drop) the PADI packet, or the **delay** action to set a delay between receipt of the PADI packet and transmission of the PADO packet.

(Optional) To configure a named service for a PPPoE service name table, do one of the following:

- Assign a service name to the table. The **terminate** action is applied to the service by default.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service gold-service
  ```

- Specify the action taken for a service in response to a PADI packet from a PPPoE client.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service gold-service delay 25
  ```

You can also accomplish the following optional tasks when you configure a named service:

- Specify the agent circuit identifier (ACI) and agent remote identifier (ARI) information to determines the action taken by the PPPoE underlying interface when it receives a PADI packet with matching ACI/ARI information.
- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Limit the number of active PPPoE sessions that the router can establish with the specified named service.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 507

- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

 Copyright © 2013, Juniper Networks, Inc.

- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 508
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information

You can configure up to 8000 agent circuit identifier/agent remote identifier (ACI/ARI) pairs per PPPoE service name table, distributed in any combination among the named, **empty**, and **any** service entries in the service name table. You can optionally configure the action taken by the PPPoE underlying interface when it receives a PADI packet that includes a service name tag and the vendor-specific tag with ACI/ARI information that matches the ACI/ARI pair that you specify.

You can use an asterisk (*) as a wildcard character to match ACI/ARI pairs, the ACI alone, or the ARI alone. The asterisk can be placed only at the beginning, the end, or both the beginning and end of the identifier string. You can also specify an asterisk alone for either the ACI or the ARI. You cannot specify only an asterisk for both the ACI and the ARI. When you specify a single asterisk as the identifier, that identifier is ignored in the PADI packet.

For example, suppose you care about matching only the ACI and do not care what value the ARI has in the PADI packet, or even whether the packet contains an ARI value. In this case you can set the ***remote-id-string*** to a single asterisk. Then the interface ignores the ARI received in the packet and the interface takes action based only on matching the specified ACI.

To indicate that it can service the client request, the interface returns a PADO packet in response to the PADI packet. By default, the interface immediately responds to the request; this is the **terminate** action. Alternatively, you can configure the **drop** action to ignore (drop) the PADI packet, or the **delay** action to set a delay between receipt of the PADI packet and transmission of the PADO packet.

To configure an ACI/ARI pair for a named, **empty**, or **any** service, do one of the following:

- Assign an ACI/ARI pair to the service name. The **terminate** action is applied to the pair by default.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service gold-service agent-specifier aci DSLAM:3/0/1/101 ari *user*
  ```

- Specify the action taken for the ACI/ARI pair in response to a PADI packet from a PPPoE client.

  ```
  [edit protocols pppoe service-name-tables table1]
  user@host# set service any agent-specifier aci velorum-ge-2/0/3 ari westford delay
    90
  ```

  In this example, an ACI/ARI pair and the **delay** action are configured for the **any** service. Configuring an ACI/ARI pair for the **any** service is useful when you want to match the agent circuit identifier and agent remote identifier information for a specific PPPoE

Copyright © 2013, Juniper Networks, Inc.

client, but do not care about the contents of the service name tag transmitted by the client in the PADI packet.

You can also accomplish the following optional tasks when you configure an ACI/ARI pair:

- Specify a dynamic profile and routing instance with which the router instantiates a dynamic PPPoE subscriber interface.
- Reserve a specified static PPPoE interface for exclusive use by the PPPoE client with match ACI/ARI information.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*
- Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client on page 509
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name

To limit the number of PPPoE client sessions that can use a particular service entry in the PPPoE service name table, you can configure the maximum number of PPPoE sessions using static or dynamic PPPoE interfaces that the router can establish with the specified named service, **empty** service, or **any** service. You cannot configure a maximum sessions limit for an ACI/ARI pair in the service name table.

The maximum sessions limit must be in the range 1 through the platform-specific maximum PPPoE sessions supported for your routing platform. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

To limit the number of PPPoE client sessions for a particular named, **empty**, or **any** service:

- Configure the maximum sessions limit for the specified service:

```
[edit protocols pppoe service-name-tables tableEast]
user@host# set service premium-service max-sessions 100
```

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- PPPoE Overview on page 484

 Copyright © 2013, Juniper Networks, Inc.

Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client

To reserve a static PPPoE interface for exclusive use by the PPPoE client with matching agent circuit identifier/agent remote identifier (ACI/ARI) information, you can assign a previously configured static PPPoE interface to an ACI/ARI pair defined for a named service entry, **empty** service entry, or **any** service entry in a PPPoE service name table. You cannot assign a static PPPoE interface directly to a service entry that does not have an ACI/ARI pair defined.

Observe the following guidelines when you configure a static PPPoE interface for an ACI/ARI pair:

- You can specify only one static PPPoE interface per ACI/ARI pair.
- If the ACI/ARI pair represents an individual PPPoE client, make sure there is a one-to-one correspondence between the client and the static PPPoE interface.
- The static interface associated with the ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.
- You cannot configure a static interface for an ACI/ARI pair already configured with a dynamic profile and routing instance. Conversely, you cannot configure a dynamic profile and routing instance for an ACI/ARI pair already configured with a static interface.

Before you begin:

- Configure the static PPPoE interface on a M120, M320, or MX Series router.

  See "Configuring PPPoE" on page 494.

To reserve a static PPPoE interface for exclusive use by the PPPoE client with matching ACI/ARI information:

- Assign a previously configured static PPPoE interface to the ACI/ARI pair defined for a named, **empty**, or **any** service entry:

```
[edit protocols pppoe service-name-tables tableEast]
user@host# set service any agent-specifier aci velorum-ge-2/0/3 ari westford
  static-interface pp0.100
```

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- PPPoE Overview on page 484

Enabling Advertisement of Named Services in PADO Control Packets

You can enable advertisement of named services in PADO control packets sent by the router to the PPPoE client to indicate the services that the router can offer. By default, advertisement of named services in PADO packets is disabled. You can enable PADO advertisement as a global option on the router when you configure the PPPoE protocol.

NOTE: **Make sure the combined number and length of all named services advertised in the PADO packet does not exceed the MTU size of the PPPoE underlying interface.**

To enable advertisement of named services in PADO packets:

- Configure the PPPoE protocol to enable PADO advertisement:

  ```
  [edit protocols pppoe]
  user@host# set pado-advertise
  ```

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- PPPoE Overview on page 484

Assigning a Service Name Table to a PPPoE Underlying Interface

You must assign the PPPoE service name table to a PPPoE underlying interface.

Before you begin:

- Specify PPPoE as the encapsulation method on the underlying interface.

  See *Setting the Appropriate Encapsulation on the PPPoE Interface* in "Configuring PPPoE" on page 494.

To assign a service name table to a PPPoE underlying interface:

- Specify the table name:

  ```
  [edit interfaces interface-name unit logical-unit-number]
  user@host# set pppoe-underlying-options service-name-table table1
  ```

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Example: Configuring a PPPoE Service Name Table on page 512
- PPPoE Overview on page 484

- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

Disabling the Sending of PPPoE Access Concentrator Tags in PADS Packets

By default, a router that functions as an access concentrator (AC) sends the AC-Name and AC-Cookie tags, along with the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags, in the PPPoE Active Discovery Session (PADS) packet when it confirms a session with a PPPoE client. The AC-Name and AC-Cookie tags are defined as follows:

- AC-Name—String that uniquely identifies the particular AC
- AC-Cookie—Tag used by the AC to help protect against denial-of-service (DoS) attacks

If it is necessary for compatibility with your network equipment, you can prevent the router from sending the AC-Name and AC-Cookie tags in the PADS packet.

To prevent the router from transmitting the AC-Name and AC-Cookie tags in the PADS messages:

- Specify that PADS messages with AC-Name and AC-Cookie tags are not sent.

  ```
  [edit protocols pppoe]
  user@host# set no-send-pads-ac-info
  ```

  The **no-send-pads-ac-info** statement affects PADS packets sent only on PPPoE interfaces configured on the router after you configure this statement. It has no effect on PADS packets sent on previously created PPPoE interfaces.

Related Documentation

- PPPoE Overview on page 484

Discarding PADR Messages to Accommodate Abnormal CPE Behavior

This topic describes how to avoid a situation where certain CPEs respond inappropriately to normal router behavior.

During PPPoE session negotiation, the router returns PADS messages in response to PADR messages when it accepts or rejects the PPPoE session. The router adds an error tag to the PADS message when it detects a problem.

AC-System-Error is one such tag. This tag is inserted when the router imposes automatic throttling in response to excessive CPU consumption, excessive subscriber connections, or physical interfaces cycling up and down.

When the CPE receives a PADS message with this tag, the typical behavior is to retry sending PADR messages to the router or to restart session negotiation by sending PADI messages. However, some CPEs may respond inappropriately with the result that their subscribers are never connected until the CPE is rebooted.

To avoid this situation when such CPEs have access to your network, you can include the **no-send-pads-error** statement at the **[edit protocols pppoe]** hierarchy level. This statement causes the router to silently discard PADR messages in situations where the PADS would include the AC-System-Error tag. The consequence is that the CPE resends

Copyright © 2013, Juniper Networks, Inc.

PADR messages. When the conditions that result in the AC-System-Error tag are no longer present, the router once again evaluates PADR packets to determine whether to accept or reject the session.

To silently discard PADR packets:

- Specify that PADS messages with AC-System-Error tags are not sent.

```
[edit protocols pppoe]
user@host# set no-send-pads-error
```

Related Documentation

- PPPoE Overview on page 484

Example: Configuring a PPPoE Service Name Table

This example shows how you can configure a PPPoE service name table on an M120 router, M320 router, or MX Series router with service entries that correspond to different client services. By configuring the appropriate actions (**delay**, **terminate**, or **drop**) and agent circuit identifier/agent remote identifier (ACI/ARI) pairs for the service entries, you can provide load balancing and redundancy across a set of remote access concentrators (ACs) in a mesh topology, and determine how best to allocate service requests from PPPoE clients to the servers in your network.

In this example, the PPPoE service name table, Table1, contains the following service entries:

- **user1-service**—Named service representing the subscriber service for user1.
- **user2-service**—Named service representing the subscriber service for user2.
- **empty** service—Represents an unspecified service.

To configure a PPPoE service name table with service entries that correspond to different subscriber services:

1. Create the PPPoE service name table and define the services and associated actions.

```
[edit protocols pppoe]
service-name-tables Table1 {
  service empty {
    drop;
  }
  service user1-service {
    terminate;
    agent-specifier {
      aci "east*" ari "wfd*" delay 10;
      aci "west*" ari "svl*" delay 10;
    }
  }
  service user2-service {
    delay 20;
  }
}
```


Copyright © 2013, Juniper Networks, Inc.

This example creates a PPPoE service name table named Table1 with three service entries, as follows:

- The **empty** service is configured with the **drop** action. This action prohibits the router (AC) from responding to PADI packets from the client that contain empty service name tags.
- The **user1-service** named service is configured with both the **terminate** action, and two ACI/ARI (agent-specifier) pairs:
  - The **terminate** action directs the router to immediately respond to PADI packets from the client that contain the **user1-service** tag, and is the default action for named services.
  - The 10-second delay configured for each ACI/ARI pair applies only to PADI packets from the client that contains a vendor-specific tag with matching ACI and ARI information. In this example, configuring the **delay** action indicates that the **east** or **west** server is considered the backup AC for handling these client requests, and that you expect an AC other than **east** or **west** to handle the request as the primary server. If the primary AC does not respond to the client with a PADO packet within 10 seconds, then the **east** or **west** backup AC sends the PADO packet after the 10-second delay expires.
- The **user2-service** named service is configured with a 20-second delay, indicating that you expect an AC other than the one on which this PPPoE service name table is configured to be the primary AC for handling this client request. If the primary AC does not respond to the client with a PADO packet within 20 seconds, then the backup AC (that is, the router on which you are configuring the service name table) sends the PADO packet after the 20-second delay expires.

2. Assign the PPPoE service name table to a PPPoE underlying interface configured with PPPoE encapsulation.

```
[edit interfaces]
ge-2/0/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    encapsulation ppp-over-ethernet;
    pppoe-underlying-options {
      service-name-table Table1;
    }
  }
}
```

3. (Optional) Verify the PPPoE service name table configuration.

```
user@host> show pppoe service-name-tables Table1
Service Name Table: Table1
  Service Name: <empty>
  Service Action: Drop

  Service Name: user1-service
  Service Action: Terminate
    ACI: east*
    ARI: wfd*
```

```
        ACI/ARI Action: Delay 10 seconds
      ACI: west*
      ARI: svl*
        ACI/ARI Action: Delay 10 seconds

    Service Name: user2-service
    Service Action: Delay 20 seconds
```

4. (Optional) Verify whether the PPPoE service name table has been properly assigned to the underlying PPPoE interface, and whether packet transfer between the router (AC) and PPPoE client is working correctly.

```
user@host>show pppoe underlying-interfaces ge-2/0/3.0 extensive
ge-2/0/3.0 Index 72
  State: Static, Dynamic Profile: None,
  Max Sessions: 4000, Active Sessions: 2,
  Service Name Table: Table1, Duplicate Protection: Off,
  AC Name: east
  PacketType                         Sent         Received
    PADI                                0                2
    PADO                                2                0
    PADR                                0                2
    PADS                                2                0
    PADT                                0                1
    Service name error                  0                0
    AC system error                     0                0
    Generic error                       0                0
    Malformed packets                   0                0
    Unknown packets                     0                0
```

Examine the command output to ensure the following:

- The **Service Name Table** field displays the name of the correct PPPoE service name table. This field displays **none** if no service name table has been associated with the specified interface.
- The **Sent** and **Received** values for the **Service name error** field are 0 (zero). For example, a nonzero value in the **Received** field for **Service name error** indicates that there are errors in the control packets received from PPPoE clients, such as a PADI packet that does not contain a service name tag.

Related Documentation

- Understanding PPPoE Service Name Tables on page 488
- Configuring PPPoE Service Name Tables on page 502
- Troubleshooting PPPoE Service Name Tables on page 516
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Tracing PPPoE Operations

The Junos OS trace feature tracks PPPoE operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

 Copyright © 2013, Juniper Networks, Inc.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

1. Important events are logged in a file called **pppoed** located in the **/var/log** directory. You cannot change the directory (**/var/log**) in which trace files are located.
2. When the file **pppoed** reaches 128 kilobytes (KB), it is renamed **pppoed.0**, then **pppoed.1**, and finally **pppoed.2**, until there are three trace files. Then the oldest trace file (**pppoed.2**) is overwritten.

   You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPPoE tracing operations:

1. Specify that you want to configure tracing options.

   ```
   [edit protocols pppoe]
   user@host# edit traceoptions
   ```
2. (Optional) Configure the name for the file used for the trace output.
3. (Optional) Configure the number and size of the log files.
4. (Optional) Configure access to the log file.
5. (Optional) Configure a regular expression to filter logging events.
6. (Optional) Configure flags to filter the operations to be logged.

Optional PPPoE traceoptions operations are described in the following sections:

Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is **pppoed**. You can specify a different name with the **file** option.

Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

Copyright © 2013, Juniper Networks, Inc.

The archived files are differentiated by a suffix in the format ***.number*.gz**. The newest archived file is **.0.gz** and the oldest archived file is ***.(maximum number)-1*.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, ***filename***, reaches 2 MB, ***filename*** is compressed and renamed ***filename*.0.gz**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0.gz** is renamed ***filename*.1.gz** and ***filename*** is compressed and renamed ***filename*.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, ***filename*.19.gz**, is simply overwritten when the next oldest file, ***filename*.18.gz** is compressed and renamed to ***filename*.19.gz**.

Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

- [edit protocols pppoe traceoptions]
  user@host# **set flag authentication**

Related Documentation

- PPPoE Overview on page 484
- *Ethernet Interfaces*

Troubleshooting PPPoE Service Name Tables

Problem A misconfiguration of a PPPoE service name table can prevent PPPoE services from being properly activated. Configuration options for PPPoE service name tables are simple, which should simplify discovering where a misconfiguration exists. PPPoE clients cannot connect if the service name table contains no match for the service name tag carried in the PADI packet.

The symptom of a service name table misconfiguration is that the client connection process stops at the negotiation stage and the PADI packets are ignored. You can use the **show pppoe statistics** command to examine the PPPoE packet counts for a problem.

 Copyright © 2013, Juniper Networks, Inc.

When the service name table is properly configured, packets sent and received increment symmetrically. The following sample output shows a PADO sent count equal to the PADI received count, and PADS sent count equal to the PADR received count. This output indicates that the PPPoE negotiation is proceeding successfully and that the service name table is not misconfigured.

```
user@host> show pppoe statistics ge-2/0/3.1

 Active PPPoE sessions: 2
    PacketType                        Sent         Received
      PADI                               0               16
      PADO                              16                0
      PADR                               0               16
      PADS                              16                0
      PADT                               0                0
      Service name error                 0                0
      AC system error                    0                0
      Generic error                      0                0
      Malformed packets                  0                0
      Unknown packets                    0                0
```

When the service name table is misconfigured, the output of the **show pppoe statistics** command indicates that the number of PADI packets received on the underlying interface is increasing, but the number of PADO packets sent remains at zero. The following sample output shows a PADI count of 100 and a PADO count of 0.

```
user@host> show pppoe statistics ge-2/0/3.1

Active PPPoE sessions: 0
    PacketType                        Sent         Received
      PADI                               0              100
      PADO                               0                0
      PADR                               0                0
      PADS                               0                0
      PADT                               0                0
      Service name error                 0                0
      AC system error                    0                0
      Generic error                      0                0
      Malformed packets                  0                0
      Unknown packets                    0                0
```

When you believe a misconfiguration exists, use the **monitor traffic interface** command on the underlying interface to determine which service name is being requested by the PPPoE client. The following sample output shows that the client is requesting Service1 in the service name tag.

```
user@host> monitor traffic interface ge-2/0/3.1 print-hex print-ascii
 Listening on ge-2/0/3.1, capture size 96 bytes

 11:49:41.436682  In PPPoE PADI [Service-Name "Service1"] [Host-Uniq UTF8]
[TAG-0x120 UTF8] [Vendor-Specific UTF8]
 0x0000   ffff ffff ffff 0090 1a42 0ac1 8100 029a     .........B......
 0x0010   8863 1109 0000 00c9 0101 0008 5365 7276     .c..........Serv
 0x0020   6963 6531 0103 0004 1200 9c43 0120 0002     ice1.......C....
 0x0030   044a 0105 00ab 0000 0de9 0124 783a 3132     .J.........$x:12
 0x0040   3030 3963                                   009c
```

You can then use the **show pppoe service-name-tables** command to determine whether you have misspelled the name of the service or perhaps not configured the service at all.

Copyright © 2013, Juniper Networks, Inc.

Cause Typical misconfigurations appear in the service name table configurations.

Solution Use the appropriate statements to correct the misconfiguration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- *show pppoe service-name-tables*
- *show pppoe statistics*
- *show pppoe underlying-interfaces*
- PPPoE Overview on page 484
- *Ethernet Interfaces*

Verifying a PPPoE Configuration

Purpose You can use show commands to display and verify the PPPoE configuration.

Action To verify a PPPoE configuration, you can issue the following operational mode commands:

- **show interfaces at-*fpc/pic/port* extensive**
- **show interfaces pp0**
- **show pppoe interfaces**
- **show pppoe version**
- **show pppoe service-name-tables**
- **show pppoe sessions**
- **show pppoe statistics**
- **show pppoe underlying-interfaces**

For more information about these operational mode commands, see *Junos OS Operational Mode Commands*.

Related Documentation

- PPPoE Overview on page 484
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 31

Configuring Ethernet Automatic Protection Switching

Ethernet Automatic Protection Switching Overview

Ethernet automatic protection switching (APS) is a linear protection scheme designed to protect VLAN based Ethernet networks.

With Ethernet APS, a protected domain is configured with two paths, a working path and a protection path. Both working and protection paths can be monitored using an Operations Administration Management (OAM) protocol like Connectivity Fault Management (CFM). Normally, traffic is carried on the working path (that is, the working path is the active path), and the protection path is disabled. If the working path fails, its protection status is marked as degraded (DG) and APS switches the traffic to the protection path, then the protection path becomes the active path.

APS uses two modes of operation, linear 1+1 protection switching architecture and linear 1:1 protection switching architecture. The linear 1+1 protection switching architecture operates with either unidirectional or bidirectional switching. The linear 1:1 protection switching architecture operates with bidirectional switching.

In the linear 1+1 protection switching architecture, the normal traffic is copied and fed to both working and protection paths with a permanent bridge at the source of the protected domain. The traffic on the working and protection transport entities is transmitted simultaneously to the sink of the protected domain, where a selection between the working and protection transport entities is made.

In the linear 1:1 protection switching architecture, the normal traffic is transported on either the working path or on the protection path using a selector bridge at the source of the protection domain. The selector at the sink of the protected domain selects the entity that carries the normal traffic.

Copyright © 2013, Juniper Networks, Inc.

Unidirectional and Bidirectional Switching

Unidirectional switching utilizes fully independent selectors at each end of the protected domain. Bidirectional switching attempts to configure the two end points with the same bridge and selector settings, even for a unidirectional failure. Unidirectional switching can protect two unidirectional failures in opposite directions on different entities.

Selective and Merging Selectors

In the linear 1:1 protection switching architecture, where traffic is sent only on the active path, there are two different ways in which the egress direction (the direction out of the protected segment) data forwarding can act: selective selectors and merging selectors. A selective selector forwards only traffic that is received from both the paths regardless of which one is currently active. In other words, with a merging selector the selection of the currently active path only affects the ingress direction. Merging selectors minimize the traffic loss during a protection switch, but they do not guarantee the delivery of the data packets in order.

Revertive and Nonrevertive Switching

For revertive switching, traffic is restored to the working path after the conditions causing the switch have cleared.

For nonrevertive switching, traffic is allowed to remain on the protection path even after the conditions causing the switch have cleared.

NOTE: **The configuration on both the provider edge (PE) routers have to be either in revertive mode or non-revertive mode.**

Protection Switching Between VPWS Pseudowires

Figure 33: Connections Terminating on Single PE

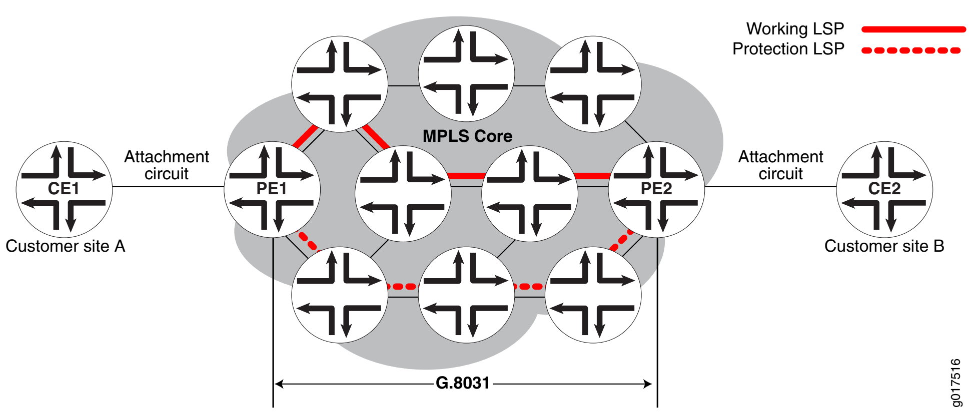

In the scenario diagramed in Figure 33 on page 520, a Virtual Private Wire Service (VPWS) is provisioned between customer sites A and B using a single pseudowire (layer 2 circuit) in the core network, and two Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) are provisioned, one for the working path and the other one for the protection

Copyright © 2013, Juniper Networks, Inc.

path. CFM CCM will be used to monitor the status of each LSP. Provider edge routers PE1 and PE2 run G.8031 Ethernet APS to select one of the LSPs as the active path. Once the active path is elected at the source end of the protection group, PE1 forwards to traffic from site A to the elected active path. At the sink end of the protection group, PE2 implements a merging selector, meaning it forwards the traffic coming from both the LSPs to the customer site B.

Figure 34: Connections Terminating on a Different PE

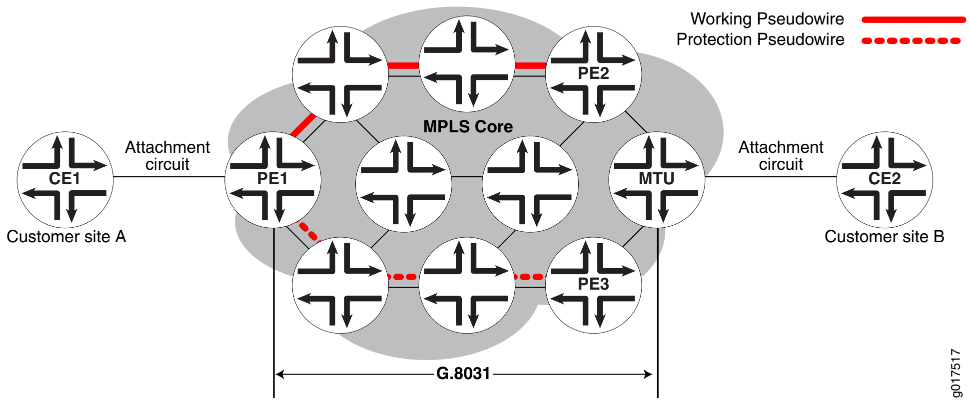

In the scenario represented in Figure 34 on page 521, a VPWS is provisioned between customer sites A and B using two pseudowires (layer 2 circuit) in the core network, one for the working path and the other for the protection path. CFM CCM will be used to monitor the status of each pseudowire.

Provider edge router PE1 and MTU run G.8031 Ethernet APS to select one of the pseudowires as the active path. Once the active path is elected at the source end of the protection group, PE1 forwards the traffic from site A to the elected active path. At the sink end of the protection group, MTU implements a merging selector, meaning it forwards the traffic coming from both the pseudowires to customer site B.

CLI Configuration Statements

```
[edit protocols protection-group]
ethernet-aps profile1{
   protocol g8031;
   revert-time seconds;
   hold-time 0-10000ms;
   local-request lockout;
}
```

revert-time- By default, protection logic restores the use of the working path once it recovers. The revert-time statement specifies how much time should elapse before the path for data should be switched from Protection to Working once recovery for Working has occurred. A revert-time of zero indicates no reversion. It will default to 300 sec (5 minutes) if not configured.

hold-time- Once a failure is detected, APS waits until this timer expires before initiating the protection switch. The range of the hold-time timer is 0 to 10,000 milliseconds. It will default to zero if not configured.

Copyright © 2013, Juniper Networks, Inc.

local-request- Configuring this value to lockout or force-switch will trigger lockout or force-switch operation on the protection groups using this profile.

Related Documentation

- Mapping of CCM Defects to APS Events on page 522
- Example: Configuring Protection Switching Between Psuedowires on page 523

Mapping of CCM Defects to APS Events

The continuity check message (CCM) engine marks the status of working and protected transport entities as either Down, Degraded, or Up.

Down—The monitored path is declared down if any of the following Multiple End Point (MEP) defects occur:

- Interface down
- CCM expiry
- RDI indicating signal failure

Degraded—The monitored path is declared degraded if any of the following MEP defects occur:

- FRR on
- FRR-ACK on

Up—The monitored path is declared up in the absence of any of the above events.

Figure 35: Understanding APS Events

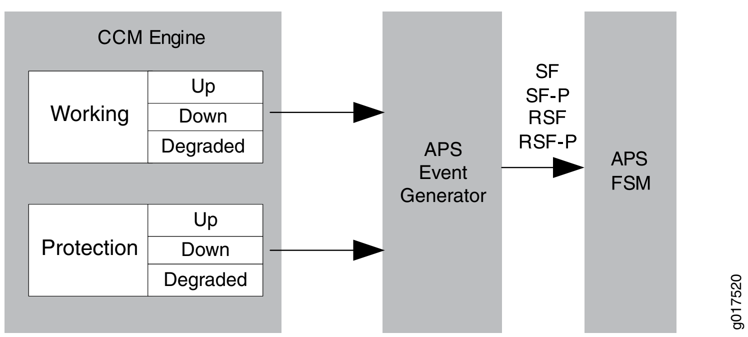

As show in Figure 35 on page 522, the APS event generator generates the following APS events based on the status of the working and protection paths:

- **SF**—Signal failure on working path
- **RSF**—Working path recovers from signal failure
- **SF-P**—Signal failure on protection path
- **RSF-P**—Protection path recovers from signal failure

Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 519
- Example: Configuring Protection Switching Between Psuedowires on page 523

Example: Configuring Protection Switching Between Psuedowires

Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.2 or later
- 2 MX Series PE routers

Overview and Topology

The physical topology of the protection switching between psuedowires example is shown in Figure 36 on page 523.

Figure 36: Topology of a Network Using VPWS Psuedowires

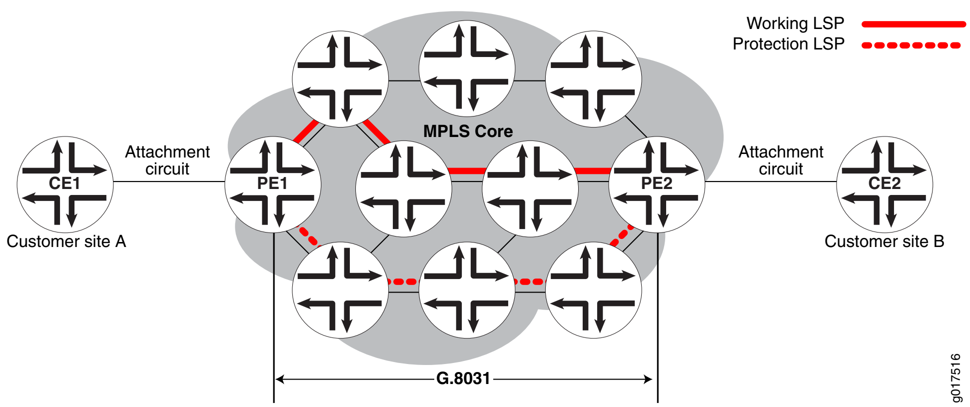

The following definitions describe the meaning of the device abbreviations used in Figure 36 on page 523.

- Customer edge (CE) device—A device at the customer site that provides access to the service provider's VPN over a data link to one or more provider edge (PE) routers.
- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.

Configuration

Step-by-Step Procedure

To configure protection switching between psuedowires, perform these tasks:

1. Configure automatic protection switching.

Copyright © 2013, Juniper Networks, Inc.

```
protocols {
   protection-group {
      ethernet-aps {
         profile-1 {
            protocol g8031;
            hold-time 1000s;
            revert-time 5m;
         }
      }
   }
```

2. Configure the connectivity fault management.

```
ethernet {
   oam {
      connectivity-fault-management {
         maintenance-domain md1 {
         level 5;
```

3. Configure the continuity check message for the working path.

```
maintenance-association W {
   protect maintenance-association P {
      aps-profile profile-1;
   }
   continuity-check {
      interval 1s;
   }
   mep 100 {
      interface ge-1/0/0.0 working;
      direction down;
      auto-discovery;
   }
}
```

4. Configure the continuity check message for the protection path.

```
maintenance-association P {
   continuity-check {
      interval 1s;
   }
   mep 100 {
      interface ge-1/0/0.0 protect;
      direction down;
      auto-discovery;
   }
}
```

Results Check the results of the configuration:

```
protocols {
   protection-group {
      ethernet-aps {
         profile-1 {
            protocol g8031;
            hold-time 1000s;
            revert-time 5m;
         }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
  }
  ethernet {
    oam {
      connectivity-fault-management {
        maintenance-domain md1 {
          level 5;
          maintenance-association W {
            protect maintenance-association P {
              aps-profile profile-1;
            }
            continuity-check {
              interval 1s;
            }
            mep 100 {
              interface ge-1/0/0.0 working;
              direction down;
              auto-discovery;
            }
          }
          maintenance-association P {
            continuity-check {
              interval 1s;
            }
            mep 100 {
              interface ge-1/0/0.0 protect;
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
```

Related Documentation

- Ethernet Automatic Protection Switching Overview on page 519
- Mapping of CCM Defects to APS Events on page 522

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 32

Configuring Ethernet Ring Protection Switching

Ethernet Ring Protection Switching Overview

Ethernet ring protection switching (ERPS) helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. The RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An automatic protection switching (APS) protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.

The following standards provide detailed information on Ethernet ring protection switching:

- IEEE 802.1Q - 1998
- IEEE 802.1D - 2004
- IEEE 802.1Q - 2003
- Draft ITU-T Recommendation G.8032/Y.1344, *Ethernet Ring protection switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet-based networks*

For additional information on configuring Ethernet ring protection switching on EX Series switches, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*.

Copyright © 2013, Juniper Networks, Inc.

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *Layer 2 Configuration Guide* for a complete example of Ethernet rings and information about STP loop avoidance and prevention.

Related Documentation

- Understanding Ethernet Ring Protection Switching Functionality on page 528
- Configuring Ethernet Ring Protection Switching on page 532
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 533

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Ethernet Interfaces*

Understanding Ethernet Ring Protection Switching Functionality

Acronyms

The following acronyms are used in the discussion about Ethernet ring protection switching:

- MA—Maintenance association
- MEP—Maintenance association end point
- OAM—Operations, administration, and management (Ethernet ring protection switching uses connectivity fault management daemon)
- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—Ring automatic protection switching
- WTR—Wait to restore
- RPL—Ring protection link

Copyright © 2013, Juniper Networks, Inc.

Ring Nodes

Multiple nodes are used to form a ring. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL. This node also initiates the RAPS message.

Ring Node States

There are three different states for each node of a specific ring:

- init—Not a participant of a specific ring.
- idle—No failure on the ring; the node is performing normally. For a normal node, traffic is unblocked on both ring ports. For the RPL owner, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- protection—A failure occurred on the ring. For a normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.

There can be only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

Failure Detection

Ethernet ring operation depends on quick and accurate failure detection. The failure condition *signal failure (SF)* is supported. For SF detection, an Ethernet continuity check MEP must be configured for each ring link. For fast protection switching, a 10-ms transmission period for this MEP group is supported. OAM monitors the MEP group's MA and reports SF or SF clear events to the Ethernet ring control module. For this MEP group, the action profile must be configured to update the interface device IFF_LINKDOWN flag. OAM updates the IFF_LINKDOWN flag to notify the Ethernet ring control module.

Logical Ring

This feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring must operate on the physical interface, not the VLAN.

FDB Flush

When ring protection switching occurs, normally an *FDB flush* is executed. The Ethernet ring control module uses the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.

Traffic Blocking and Forwarding

Ethernet ring control uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical

Copyright © 2013, Juniper Networks, Inc.

interface default STP index state to forwarding or discarding in order to control user traffic.

RAPS Message Blocking and Forwarding

The router or switch treats the ring automatic protection switching (RAPS) message the same as it treats user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS messages between the two ring ports, as shown in Figure 37 on page 530, the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces' STP index state. The RAPS message is always sent by the router or switch through the ring ports, as shown in Figure 38 on page 530. A RAPS message received from a discarding ring port is sent to the Ethernet ring control module, but is not sent to the other ring port.

Figure 37: Protocol Packets from the Network to the Router

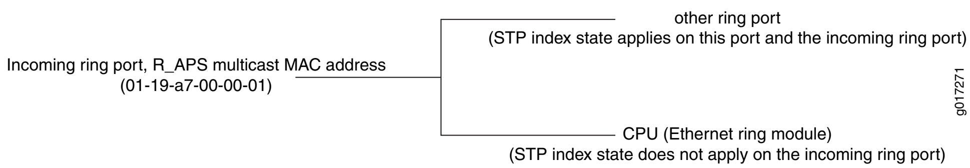

Figure 38: Protocol Packets from the Router or Switch to the Network

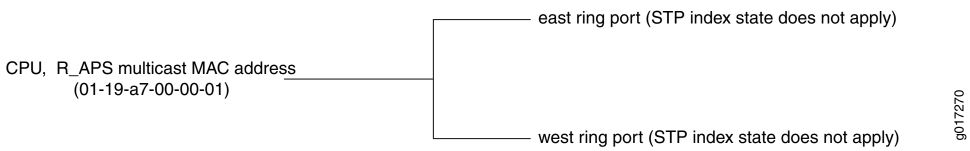

Juniper Networks EX Series switches and Juniper Networks MX Series routers use different methods to achieve these routes.

The switches use forwarding database entries to direct the RAPS messages. The forwarding database entry (keyed by the RAPS multicast address and VLAN) has a composite next hop associated with it—the composite next hop associates the two ring interfaces with the forwarding database entry and uses the split horizon feature to prevent sending the packet out on the interface that it is received on. This is an example of the forwarding database entry relating to the RAPS multicast MAC (a result of the **show ethernet-switching table detail** command):

```
VLAN: v1, Tag: 101, MAC: 01:19:a7:00:00:01, Interface: ERP
Interfaces:          ge-0/0/9.0, ge-0/0/3.0
Type: Static
Action: Mirror
Nexthop index: 1333
```

The routers use an implicit filter to achieve ERP routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain

Copyright © 2013, Juniper Networks, Inc.

the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:
  - term 1: if [Ethernet type is not OAM Ethernet type (0x8902)
    ]   { accept packet }
  - term 2: if [source MAC address belongs to this bridge]
          { drop packet, our packet loop through the ring and come back
           to home}
  - term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,
           0x00,0x00,0x01] AND[ring port STP status is DISCARDING]
          { send to CPU }
- Control channel related terms:
  - if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,
     0x01] AND[ring port STP status is FORWARDING] AND [Incoming interface
     IFL equal to control channel IFL]
    { send packet to CPU and send to the other ring port }
    default term: accept packet.

Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain for routers (or the same VLAN for switches) in order to forward RAPS protocol data units (PDUs) between the two ring control physical interfaces. If the router control channel logical interface is not a trunk port, only control logical interfaces will be configured in ring port configuration. If this router control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured for routers. For EX Series switches, always specify either a VLAN name or VLAN ID for all links.

RAPS Message Termination

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). However, interconnection of multiple rings is not supported in this release. The interconnection of two rings means that two rings may share the same link or share the same node.

Node ID

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address.

For routers only, you can configure this node ID when configuring the ring on the node or automatically select an ID such as STP. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID. The node ID on EX Series switches is selected automatically and is not configurable.

Bridge Domains with the Ring Port (MX Series Routers Only)

On the routers, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain, you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB Layer 3 interface.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527
- Configuring Ethernet Ring Protection Switching on page 532
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 533

- *Ethernet Interfaces*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Configuring Ethernet Ring Protection Switching

The inheritance model follows:

```
protection-group {
  ethernet-ring ring-name (
    node-id mac-address;
    ring-protection-link-owner;
    east-interface {
      control-channel channel-name {
      ring-protection-link-end;
    }
    west-interface {
      node-id mac-address;
      control-channel channel-name {
      ring-protection-link-end;
    }
    data-channel {
      vlan number;
    }
    guard-interval number;
```


Copyright © 2013, Juniper Networks, Inc.

```
        restore-interval number;
      }
    }
```

For each ring, a protection group must be configured. There may be several rings in each node, so there should be multiple protection groups corresponding to the related Ethernet rings.

Three interval parameters (**restore-interval**, **guard-interval**, and **hold-interval**) can be configured at the protection group level. These configurations are global configurations and apply to all Ethernet rings if the Ethernet ring doesn't have a more specific configuration for these values. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527
- Understanding Ethernet Ring Protection Switching Functionality on page 528
- Example: Ethernet Ring Protection Switching Configuration on MX Routers on page 533

- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Ethernet Interfaces*

Example: Ethernet Ring Protection Switching Configuration on MX Routers

This example describes how to configure Ethernet ring protection switching on an MX Series router:

Requirements

This example uses the following hardware and software components:

- Router node 1 running Junos OS with two Gigabit Ethernet interfaces.
- Router node 2 running Junos OS with two Gigabit Ethernet interfaces.
- Router node 3 running Junos OS with two Gigabit Ethernet interfaces.

Ethernet Ring Overview and Topology

This section describes a configuration example for a three-node ring. The ring topology is shown in Figure 39 on page 534.

Figure 39: Example of a Three-Node Ring Topology

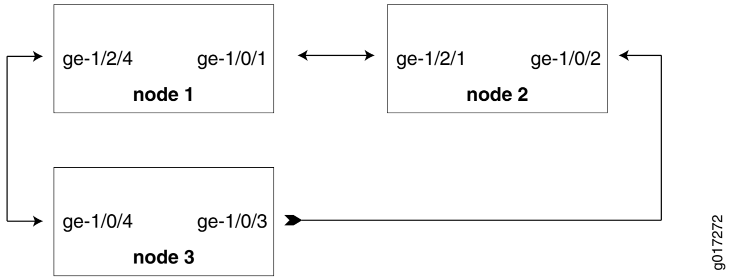

The configuration in this section is only for the RAPS channel. The bridge domain for user traffic is the same as the normal bridge domain. The only exception is if a bridge domain includes a ring port, then it must also include the other ring port of the same ring.

Configuring a Three-Node Ring

To configure Ethernet Ring Protection Switching on a three-node ring, perform these tasks:

- Configuring Ethernet Ring Protection Switching on a Three-Node Ring on page 534

Configuring Ethernet Ring Protection Switching on a Three-Node Ring

Step-by-Step Procedure

1. **Configuring Node 1**

```
interfaces {
  ge-1/0/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  ge-1/2/4 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
      }
    }
  }
  bridge-domains {
    bd1 {
      domain-type bridge;
      interface ge-1/2/4.1;
      interface ge-1/0/1.1;
      }
    }
  protocols {
    protection-group {
      ethernet-ring pg101 {
        node-id 00:01:01:00:00:01;
```

 Copyright © 2013, Juniper Networks, Inc.

```
                ring-protection-link-owner;
                east-interface {
                    control-channel ge-1/0/1.1;
                    ring-protection-link-end;
                    }
                west-interface {
                    control-channel ge-1/2/4.1;
                    }
                }
            }
        }
    protocols {
        oam {
            ethernet {
                connectivity-fault-management {
                    action-profile rmep-defaults {
                        default-action {
                            interface-down;
                            }
                        }
                    maintenance-domain d1 {
                        level 0;
                        maintenance-association 100 {
                            mep 1 {
                                interface ge-1/0/1;
                                remote-mep 2 {
                                    action-profile rmep-defaults;
                                    }
                                }
                            }
                        }
                    maintenance-domain d2 {
                        level 0;
                        maintenance-association 100 {
                            mep 1 {
                                interface ge-1/2/4;
                                remote-mep 2 {
                                    action-profile rmep-defaults;
                                    }
                                }
                            }
                        }
                    }
                }
            }
        }
    }
```

2. **Configuring Node 2**

```
interfaces {
    ge-1/0/2 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
```

```
    }
  }

  ge-1/2/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
}

bridge-domains {
  bd1 {
    domain-type bridge;
    interface ge-1/2/1.1;
    interface ge-1/0/2.1;
  }
}


protocols {
  protection-group {
    ethernet-ring pg102 {
      east-interface {
        control-channel ge-1/0/2.1;
      }
      west-interface {
        control-channel ge-1/2/1.1;
      }
    }
  }
}


protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile rmep-defaults {
          default-action {
            interface-down;
          }
        }
        maintenance-domain d1 {
          level 0;
          maintenance-association 100 {
            mep 2 {
              interface ge-1/2/1;
              remote-mep 1 {
                action-profile rmep-defaults;
              }
            }
          }
        }
```

 Copyright © 2013, Juniper Networks, Inc.

```
                maintenance-domain d3 {
                    level 0;
                    maintenance-association 100 {
                        mep 1 {
                            interface ge-1/0/2;
                            remote-mep 2 {
                                action-profile rmep-defaults;
                            }
                        }
                    }
                }
            }
        }
    }
}
```

3. **Configuring Node 3**

```
interfaces {
    ge-1/0/4 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }

    ge-1/0/3 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
}

bridge-domains {
    bd1 {
        domain-type bridge;
        interface ge-1/0/4.1;
        interface ge-1/0/3.1;
    }
}


protocols {
    protection-group {
        ethernet-ring pg103 {
            east-interface {
                control-channel ge-1/0/3.1;
            }
            west-interface {
                control-channel ge-1/0/4.1;
            }
        }
```

```
    }
  }

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile rmep-defaults {
          default-action {
            interface-down;
          }
        }
        maintenance-domain d2 {
          level 0;
          maintenance-association 100 {
            mep 2 {
              interface ge-1/0/4;
              remote-mep 1 {
                action-profile rmep-defaults;
              }
            }
          }
        }
        maintenance-domain d3 {
          level 0;
          maintenance-association 100 {
            mep 2 {
              interface ge-1/0/3;
              remote-mep 1 {
                action-profile rmep-defaults;
              }
            }
          }
        }
      }
    }
  }
}
```

Examples: Ethernet RPS Output

This section provides output examples based on the configuration shown in "Example: Ethernet Ring Protection Switching Configuration on MX Routers" on page 533. The show commands used in these examples can help verify configuration and correct operation.

Normal Situation—RPL Owner Node

If the ring has no failure, the **show** command will have the following output for Node 1:

user@node1> **show protection-group ethernet-ring aps**

```
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               NR             No        Yes

Originator  Remote Node ID
Yes
```

user@node1> **show protection-group ethernet-ring interface**

 Copyright © 2013, Juniper Networks, Inc.

```
Ethernet ring port parameters for protection group pg101

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/0/1     ge-1/0/1.1       discarding     Yes
ge-1/2/4     ge-1/2/4.1       forwarding     No

Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```

```
user@node1> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event        Ring Protection Link Owner
pg101            idle         NR-RB        Yes

Restore Timer  Quard Timer  Operation state
disabled       disabled    operational
```

```
user@node1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                          : 1
RAPS received                      : 0
Local SF happened:                 : 0
Remote SF happened:                : 0
NR event happened:                 : 0
NR-RB event happened:              : 1
```

Normal Situation—Other Nodes

For Node 2 and Node 3, the outputs should be the same:

```
user@node2> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               NR             No        Yes

Originator  Remote Node ID
No          00:01:01:00:00:01
```

```
user@node2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/2/1     ge-1/2/1.1       forwarding     No
ge-1/0/2     ge-1/0/2.1       forwarding     No

Signal Failure  Admin State
Clear          IFF ready
Clear          IFF ready
```

```
user@node2> show protection-group ethernet-ring node-state
Ethernet ring    APS State    Event        Ring Protection Link Owner
pg102            idle         NR-RB        No

Restore Timer  Quard Timer  Operation state
disabled       disabled    operational
```

```
user@node2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg101
RAPS sent                          : 0
RAPS received                      : 1
Local SF happened:                 : 0
Remote SF happened:                : 0
```

```
NR event happened:              : 0
NR-RB event happened:           : 1
```

Failure Situation—RPL Owner Node

If the ring has a link failure between Node 2 and Node 3, the **show** command will have the following outputs for Node 1:

```
user@node1> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg101               SF             NO        No

Originator  Remote Node ID
No          00:01:02:00:00:01
```

```
user@node1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/0/1     ge-1/0/1.1       forwarding     Yes
ge-1/2/4     ge-1/2/4.1       forwarding     No

Signal Failure  Admin State
Clear           IFF ready
Clear           IFF ready
```

```
user@node1> show protection-group ethernet-ring node-state
Ethernet ring     APS State    Event         Ring Protection Link Owner
pg101             protected    SF            Yes

Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

```
user@node1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent                       : 1
RAPS received                   : 1
Local SF happened:              : 0
Remote SF happened:             : 1
NR event happened:              : 0
NR-RB event happened:           : 1
```

Failure Situation—Other Nodes

For Node 2 and Node 3, the outputs should be the same:

```
user@node2> show protection-group ethernet-ring aps
Ethernet Ring Name  Request/state  No Flush  Ring Protection Link Blocked
pg102               SF             No        No

Originator  Remote Node ID
Yes         00:00:00:00:00:00
```

```
user@node2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel  Forward State  Ring Protection Link End
ge-1/2/1     ge-1/2/1.1       forwarding     No
ge-1/0/2     ge-1/0/2.1       discarding     No

Signal Failure  Admin State
```

 Copyright © 2013, Juniper Networks, Inc.

```
Clear           IFF ready
set             IFF ready
```

user@node2> **show protection-group ethernet-ring node-state**

```
Ethernet ring    APS State    Event        Ring Protection Link Owner
pg102            idle         NR-RB        No

Restore Timer  Quard Timer  Operation state
disabled       disabled     operational
```

user@node2> **show protection-group ethernet-ring statistics group-name pg102**

```
Ethernet Ring statistics for PG pg101
RAPS sent                         : 1
RAPS received                     : 1
Local SF happened:                : 1
Remote SF happened:               : 0
NR event happened:                : 0
NR-RB event happened:             : 1
```

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527
- Understanding Ethernet Ring Protection Switching Functionality on page 528
- Configuring Ethernet Ring Protection Switching on page 532
- *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 33

Example Ethernet Configurations

Example: Configuring Fast Ethernet Interfaces

The following configuration is sufficient to get a Fast Ethernet interface up and running. By default, IPv4 Fast Ethernet interfaces use Ethernet version 2 encapsulation.

```
[edit]
user@host# set interfaces fe-5/2/1 unit 0 family inet address local-address
user@host# show
interfaces {
  fe-5/2/1 {
    unit 0 {
      family inet {
        address local-address;
      }
    }
  }
}
```

Related Documentation

- *Ethernet Interfaces*

Example: Configuring Gigabit Ethernet Interfaces

The following configuration is sufficient to get a Gigabit Ethernet, Tri-Rate Ethernet copper, or 10-Gigabit Ethernet interface up and running. By default, IPv4 Gigabit Ethernet interfaces on MX Series, M Series, and T Series routers use 802.3 encapsulation. J Series Gigabit Ethernet interfaces do not support 802.3 encapsulation.

```
[edit]
user@host# set interfaces ge-2/0/1 unit 0 family inet address local-address
user@host# show
interfaces {
  ge-2/0/1 {
    unit 0 {
```

Copyright © 2013, Juniper Networks, Inc.

```
            family inet {
                address local-address;
            }
        }
    }
}
```

The M160, M320, M120, T320, and T640 2-port Gigabit Ethernet PIC supports two independent Gigabit Ethernet links.

Each of the two interfaces on the PIC is named:

ge-*fpc*/*pic*/[0.1]

Each of these interfaces has functionality identical to the Gigabit Ethernet interface supported on the single-port PIC.

Related Documentation

- *Ethernet Interfaces*

Example: Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
aggregated-devices {
    ethernet {
        device-count 15;
    }
}
```

```
[edit interfaces]
ge-1/3/0 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/0/1 {
    gigether-options {
        802.3ad ae0;
    }
}
ae0 {
    aggregated-ether-options {
        link-speed 1g;
        minimum-links 1;
    }
    }
    vlan-tagging;
    unit 0 {
        vlan-id 1;
        family inet {
            address 14.0.100.50/24;
        }
```

 Copyright © 2013, Juniper Networks, Inc.

```
    }
    unit 1 {
        vlan-id 1024;
        family inet {
            address 14.0.101.50/24;
        }
    }
    unit 2 {
        vlan-id 1025;
        family inet {
            address 14.0.102.50/24;
        }
    }
    unit 3 {
        vlan-id 4094;
        family inet {
            address 14.0.103.50/24;
        }
    }
}
```

Related Documentation

- *Ethernet Interfaces*
- Configure 'link-speed' for Gigabit Ethernet based Aggregate Ethernet interface bundles

Example: Configuring Aggregated Ethernet Link Protection

The following configuration enables link protection on the **ae0** interface, and specifies the **ge-1/0/0** interface as the primary link and **ge-1/0/1** as the secondary link.

```
[edit interfaces]
ae0 {
    aggregated-ether-options {
        link protection;
    }
}
[edit interfaces]
ge-1/0/0 {
    gigether-options {
        802.3ad ae0 primary;
    }
}
[edit interfaces]
ge-1/0/1 {
    gigether-options {
        802.3ad ae0 backup;
    }
}
```

Related Documentation

- *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

PART 3

Ethernet Interface Configuration Statements

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 34

Summary of Ethernet Interfaces Configuration Statements

The following descriptions explain each of the interface configuration statements. The statements are organized alphabetically.

802.3ad

Syntax

```
802.3ad {
  ae interface-number (primary | backup);
  lacp {
    port-priority;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced before Junos OS Release 7.4.
primary and **backup** options added in Junos OS Release 8.3.

Description Specify aggregated Ethernet logical interface number.

Options **ae** ***interface-number***—Aggregated Ethernet logical interface number.

Range: 0 through 15

primary—For link protection configurations, specify the primary link for egress traffic.

backup—For link protection configurations, specify the backup link for egress traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an Aggregated Ethernet Interface on page 87
- Configuring Aggregated Ethernet Link Protection on page 174

account-layer2-overhead (Interface Level)

Syntax

```
account-layer2-overhead {
  value;
  egress bytes;
  ingress bytes;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Add the configured number of bytes, which is the octet adjustment per packet , to the total octet count for both the ingress and egress traffic on the logical interface, or separately for ingress and egress traffic on the logical interface.

Options ***bytes***—Number of bytes added to a packet entering and exiting an interface.

Range: –128–+127 bytes

Default: **0**

egress—Include the Layer 2 overhead value in bytes that is the octet adjustment per packet in the total octet count for egress traffic only on the logical interface

ingress—Include the Layer 2 overhead value in bytes that is the octet adjustment per packet in the total octet count for ingress traffic only on the logical interface

value—Layer 2 overhead value in bytes that is the octet adjustment per packet in the total octet count for both ingress and egress traffic on the logical interface

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics on page 407
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 410
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 412
- **account-layer2-overhead (PIC Level) on page 551**
- [edit interfaces] Hierarchy Level on page 3

Copyright © 2013, Juniper Networks, Inc.

account-layer2-overhead (PIC Level)

Syntax account-layer2-overhead;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced in Junos OS Release 13.2.

Description Enable the automatic adjustment of Layer 2 overhead in bytes, which is the octet adjustment per packet, based on the encapsulation on the logical interface for the total octet count for ingress and egress traffic on all the interfaces in the PIC.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics on page 407
- Configuring Layer 2 Overhead Accounting in Interface Statistics on page 410
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics on page 412
- **account-layer2-overhead (Interface Level) on page 550**
- *[edit chassis] Hierarchy Level*

advertisement-interval

Syntax advertisement-interval *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6 for MX Series and T Series routers.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description For MX Series and T Series routers and EX Series switches, configure an interval for LLDP advertisement.

For switches configured for Link Layer Discovery Protocol, configure the frequency at which LLDP advertisements are sent.

The **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value, or an error will be returned when you attempt to commit the configuration.

NOTE: The default value of **transmit-delay** is 2 seconds. If you configure the **advertisement-interval** as less than 8 seconds and you do not configure a value for **transmit-delay**, the default value of **transmit-delay** is automatically changed to 1 second in order to satisfy the requirement that the **advertisement-interval** value must be greater than or equal to four times the **transmit-delay** value.

Default Disabled.

Options *seconds*—Interval between LLDP advertisement.

Default: 30

Range: 5 through 32768

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226
- *show lldp*
- *Configuring LLDP (CLI Procedure)*
- *Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches*
- *transmit-delay*
- *Understanding LLDP*

 Copyright © 2013, Juniper Networks, Inc.

agent-specifier

Syntax

```
agent-specifier {
  aci circuit-id-string ari remote-id-string {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    static-interface interface-name;
  }
}
```

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*]

Release Information Statement introduced in Junos OS Release 10.0.
drop, **delay**, **terminate**, **dynamic-profile**, **routing-instance**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Specify the action taken by the interface for the specified agent circuit identifier/agent remote identifier (ACI/ARI) pair when the interface receives a PPPoE Active Discovery Initiation (PADI) control packet that includes the vendor-specific tag with ACI/ARI pair information. You can configure an ACI/ARI pair for a named service, **empty** service, or **any** service in a PPPoE service name table. A maximum of 8000 ACI/ARI pairs are supported per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the named, **empty**, and **any** service entries in the service name table.

You can use an asterisk (*) as a wildcard character to match ACI/ARI pairs, the ACI alone, or the ARI alone. The asterisk can be placed only at the beginning, the end, or both the beginning and end of the identifier string. You can also specify an asterisk alone for either the ACI or the ARI. You cannot specify only an asterisk for both the ACI and the ARI. When you specify a single asterisk as the identifier, that identifier is ignored in the PADI packet.

For example, suppose you care about matching only the ACI and do not care what value the ARI has in the PADI packet, or even whether the packet contains an ARI value. In this case you can set the ***remote-id-string*** to a single asterisk. Then the interface ignores the ARI received in the packet and the interface takes action based only on matching the specified ACI.

Default The default action is terminate.

Options aci *circuit-id-string*—Identifier for the agent circuit ID that corresponds to the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

ari *remote-id-string*—Identifier for the subscriber associated with the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

The remaining statements are explained separately.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring PPPoE Service Name Tables on page 502
- Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information on page 507

aggregate (Gigabit Ethernet CoS Policer)

Syntax
```
aggregate {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a policer to apply to nonpremium traffic.

The statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Gigabit Ethernet Policers on page 397
- premium (Hierarchical Policer) on page 630
- ieee802.1p on page 579

 Copyright © 2013, Juniper Networks, Inc.

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  ethernet-switch-profile {
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
    (mac-learn-enable | no-mac-learn-enable);
  }
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    link-protection {
      disable;
    (revertive | non-revertive);
    periodic interval;
    system-priority priority;
    system-id system-id;
  }
  link-protection;
  link-speed speed;
  logical-interface-chassis-redundancy;
  logical-interface-fpc-redundancy;
  (loopback | no-loopback);
  minimum-links number;
  rebalance-periodic time hour:minute <interval hours>;
  source-address-filter {
    mac-address;
    (source-filtering | no-source-filtering);
  }
}
```

Hierarchy Level [edit interfaces ae*x*]

Copyright © 2013, Juniper Networks, Inc.

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure aggregated Ethernet-specific interface properties.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Interfaces Overview on page 35

 Copyright © 2013, Juniper Networks, Inc.

auto-negotiation

Syntax (auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-offline)>;

Hierarchy Level [edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 8.4 for J Series Services Routers.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. For EX Series switches and J Series Services Routers, explicitly enable autonegotiation only.

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On J Series Services Routers with universal Physical Interface Modules (uPIMs) and on EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

NOTE: On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

Copyright © 2013, Juniper Networks, Inc.

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options **remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

Default: local-interface-online

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Gigabit Ethernet Autonegotiation Overview on page 415
- Configuring J Series Services Router Switching Interfaces on page 40
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

Syntax bandwidth-limit *bps*;

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* aggregate],
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* premium]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define a policer to apply to nonpremium traffic.

Options *bps*—Bandwidth limit, in bits per second. Specify either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 32 Kbps through 32 gigabits per second (Gbps). For IQ2 and IQ2-E interfaces 65,536 bps through 1 Gbps. For 10-Gigabit IQ2 and IQ2-E interfaces 65,536 bps through 10 Gbps.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397
- **burst-size-limit (Policer for Gigabit Ethernet Interfaces) on page 559**

 Copyright © 2013, Juniper Networks, Inc.

burst-size-limit (Policer for Gigabit Ethernet Interfaces)

Syntax

burst-size-limit *bytes*;

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* aggregate],
[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name* premium]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define a policer to apply to nonpremium traffic.

Options

bytes—Burst length.

Range: 1500 through 100,000,000 bytes

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397
- **bandwidth-limit (Policer for Gigabit Ethernet Interfaces) on page 558**

classifier

Syntax

```
classifier {
  per-unit-scheduler {
    forwarding-class class-name {
      loss-priority (high | low);
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the classifier for the output priority map to be applied to outgoing frames on this interface.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Output Priority Map on page 399
- **input-priority-map on page 583**

Copyright © 2013, Juniper Networks, Inc.

delay (PPPoE Service Name Tables)

Syntax delay *seconds*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description Configure the PPPoE underlying interface on the router to wait a specified number of seconds after receiving a PPPoE Active Discovery Initiation (PADI) control packet from a PPPoE client before sending a PPPoE Active Discovery Offer (PADO) packet to indicate that it can service the client request

The router (PPPoE server) does not check whether another server has already sent a PADO packet during the delay period in response to the PPPoE client's PADI packet. It is up to the PPPoE client to determine whether another PPPoE server has responded to its PADI request, or if it must respond to the delayed PADO packet to establish a PPPoE session.

Options ***seconds***—Number of seconds that the PPPoE underlying interface waits after receiving a PADI packet from a PPPoE client before sending a PADO packet in response.

Range: 1 through 120 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502

 Copyright © 2013, Juniper Networks, Inc.

disable

Syntax disable;

Hierarchy Level [edit protocols lldp],
[edit protocols lldp interface (all | *interface-name*)],
[edit routing-instances *routing-instance-name* protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp interface (all | *interface-name*)]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX Series switches) Disable LLDP globally or on an interface.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226

drop (PPPoE Service Name Tables)

Syntax drop;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description Direct the router to drop (ignore) a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client that contains the specified service name tag or agent circuit identifier/agent remote identifier (ACI/ARI) information. This action effectively denies the client's request to provide the specified service, or to accept requests from the subscriber or subscribers represented by the ACI/ARI information.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502

Copyright © 2013, Juniper Networks, Inc.

dynamic-profile (PPPoE Service Name Tables)

Syntax dynamic-profile *profile-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Specify a dynamic profile to instantiate a dynamic PPPoE interface. You can associate a dynamic profile with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The dynamic profile associated with a service entry in a PPPoE service name table overrides the dynamic profile associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **dynamic-profile** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **dynamic-profile** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options ***profile-name***—Name of the dynamic profile that the router uses to instantiate a dynamic PPPoE interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring PPPoE Service Name Tables on page 502
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

 Copyright © 2013, Juniper Networks, Inc.

ethernet (Protocols OAM)

Syntax

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
    }
    performance-monitoring {
      delegate-server-processing;
      hardware-assisted-timestamping;
      sla-iterator-profiles {
        profile-name {
          disable;
          calculation-weight {
            delay delay-weight;
            delay-variation delay-variation-weight;
          }
          cycle-time milliseconds;
          iteration-period connections;
          measurement-type (loss | statistical-frame-loss | two-way-delay);
        }
      }
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      name-format (character-string | none | dns | mac+2octet);
      maintenance-association ma-name {
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
        protect-maintenance-association protect-ma-name;
        remote-maintenance-association remote-ma-name;
        continuity-check {
          convey-loss-threshold;
          hold-interval minutes;
          interface-status-tlv;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
          port-status-tlv;
        }
        mep mep-id {
          auto-discovery;
          direction (up | down);
          interface interface-name (protect | working);
          lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
            rem-err-xcon | xcon );
          priority number;
          remote-mep mep-id {
            action-profile profile-name;
            sla-iterator-profile profile-name {
```

```
                data-tlv-size size;
                iteration-count count-value;
                priority priority-value;
              }
            }
          }
        }
      }
    }
    evcs evc-id {
      evc-protocol cfm management-domain domain-id (management-association
        association-id | vpls (routing-instance instance-id);
      remote-uni-count count;
      multipoint-to-multipoint;
    }
    link-fault-management {
      action-profile profile-name {
        action {
          link-down;
          send-critical-event;
          syslog;
        }
        event {
          link-adjacency-loss;
          link-event-rate {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
          protocol-down;
        }
      }
      interface interface-name {
        apply-action-profile;
        link-discovery (active | passive);
        pdu-interval interval;
        pdu-threshold threshold-value;
        remote-loopback;
        event-thresholds {
          frame-error count;
          frame-period count;
          frame-period-summary count;
          symbol-period count;
        }
        negotiation-options {
          allow-remote-loopback;
          no-allow-link-events;
        }
      }
    }
    lmi {
      status-counter count;
      polling-verification-timer value;
      interface name {
        uni-id uni-name;
```

 Copyright © 2013, Juniper Networks, Inc.

```
      status-counter number;
      polling-verification-timer value;
      evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
      evc evc-name {
        default-evc;
        vlan-list vlan-id-list;
      }
    }
  }
}
```

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches, and M320, MX Series, and T Series routers, provide fault signaling and detection for 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 374
- *Example: Configuring Connectivity Fault Management for a PBB Network on MX Series Routers*

ethernet-policer-profile

Syntax

```
ethernet-policer-profile {
  input-priority-map {
    ieee802.1p premium [ values ];
  }
  output-priority-map {
    classifier {
      premium {
        forwarding-class class-name {
          loss-priority (high | low);
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

NOTE: On QFX Series standalone switches, this statement hierarchy is only supported on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure a class of service (CoS)-based policer. Policing applies to the inner VLAN identifiers, not to the outer tag. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the **premium** policer is not supported.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397

 Copyright © 2013, Juniper Networks, Inc.

ethernet-ring

Syntax

```
ethernet-ring ring-name {
  control-vlan (vlan-id | vlan-name);
  data-channel {
    vlan number
  }
  east-interface {
    control-channel channel-name {
      vlan number;
    }
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
  west-interface {
    control-channel channel-name {
      vlan number;
    }
  }
}
```

Hierarchy Level [edit protocols protection-group]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

Options *ring-name*—Name of the Ethernet protection ring.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

Copyright © 2013, Juniper Networks, Inc.

ethernet-switch-profile

Syntax

```
ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [ values ];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
    policer cos-policer-name {
      aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
    }
    tag-protocol-id tpid;
  }
  (mac-learn-enable | no-mac-learn-enable);
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* aggregated-ether-options]
[edit interfaces *interface-name* ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

NOTE: On QFX Series standalone switches, the **ethernet-policer-profile** CLI hierarchy and the **mac-learn-enable** statement are only supported on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure VLAN tag and MAC address accounting and filtering properties.

The statements are explained separately.

 Copyright © 2013, Juniper Networks, Inc.

NOTE: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*]** hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

Default If the **ethernet-switch-profile** statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397
- Configuring MAC Address Filtering on page 401
- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 197

fast-aps-switch

Syntax fast-aps-switch;

Hierarchy Level [edit interfaces *interface-name* 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

- *Reducing APS Switchover Time in Layer 2 Circuits*

 Copyright © 2013, Juniper Networks, Inc.

fastether-options

Syntax

```
fastether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (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;
  }
  (source-filtering | no-source-filtering);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Fast Ethernet-specific interface properties.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Interfaces Overview on page 35

fnp

Syntax

```
fnp {
  interval <100ms | 1s | 10s | 1m | 10m>;;
  loss-threshold number
  interface interface name {
    domain-id domain-id
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Command introduced in Junos OS Release 11.4.

Description On routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet failure notification protocol.

Options **interval *number***—Specifies the time between the transmission of FNP messages.

loss-threshold *number*—FNP messages that can be lost before the FNP message is considered aged out and flushed.

interface *interface-name*—Name of the Ethernet interface.

domain-id *number*—Domain ID of the access network.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Failure Notification Protocol Overview on page 314
- Configuring the Failure Notification Protocol on page 365

 Copyright © 2013, Juniper Networks, Inc.

flow-control

Syntax

(flow-control | no-flow-control);

Hierarchy Level

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* multiservice-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description

For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Default

Flow control is enabled.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Flow Control on page 45
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

flow-control-options

Syntax

```
flow-control-options {
  down-on-flow-control;
  dump-on-flow-control;
  reset-on-flow-control;
}
```

Hierarchy Level [edit interfaces mo-*fpc/pic/port* multiservice-options]

Release Information Statement introduced before Junos OS Release 8.4.

Description Configure the flow control options for application recovery in case of a prolonged flow control failure.

- **down-on-flow-control**—Bring interface down during prolonged flow control.
- **dump-on-flow-control**—Cause core dump during prolonged flow control.
- **reset-on-flow-control**—Reset interface during prolonged flow control.

Usage Guidelines See *Configuring Flow Monitoring*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

forwarding-class (Gigabit Ethernet IQ Classifier)

Syntax

```
forwarding-class class-name {
  loss-priority (high | low);
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ interfaces only, define forwarding class name and option values.

Options ***class-name***—Name of forwarding class.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Output Priority Map on page 399
- input-priority-map on page 583
- **forwarding-class** statement in the *Junos OS Class of Service Library for Routing Devices*

 Copyright © 2013, Juniper Networks, Inc.

forwarding-mode (100-Gigabit Ethernet)

Syntax

```
forwarding-mode {
  (sa-multicast | ...the following vlan-steering statement...);
  vlan-steering {
    vlan-rule (high-low | odd-even);
  }
}
```

Hierarchy Level [edit chassis fpc *slot* pic *slot*]

Release Information Statement introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1 for MX Series routers.

Description Configure the interoperation mode for 100-Gigabit Ethernet PIC or the 100-Gigabit Ethernet MIC.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 460
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4)(Type 4 1X100GE PIC for STFPC4 FPC) Using SA Multicast Mode*
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 464
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 465
- **sa-multicast (100-Gigabit Ethernet) on page 640**
- **vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 683**
- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 684**

Copyright © 2013, Juniper Networks, Inc.

framing (10-Gigabit Ethernet Interfaces)

Syntax framing (lan-phy | wan-phy);

Hierarchy Level [edit interfaces xe-*fpc/pic/port*]

[edit interfaces et-*fpc/pic/port*] (PTX Series Packet Transport Routers)

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 12.3R2 for PTX Series Packet Transport Routers.

Description For routers supporting the 10-Gigabit Ethernet interface, configure the framing format. WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600,T4000, and PTX Series Packet Transport Routers routers only.

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- On PTX Series routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ .
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.

Default Operates in LAN PHY mode.

Options **lan-phy**—10GBASE-R interface framing format that bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface.

wan-phy—10GBASE-W interface framing format that allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and SONET devices.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- 10-Gigabit Ethernet Framing Overview on page 443
- *Configuring SONET Options for 10-Gigabit Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

gigether-options

Syntax

```
gigether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (asynchronous-notification | no-asynchronous-notification);
  (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
    local-interface-offline>;
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  no-auto-mdix
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
  speed
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [ tpids ];
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [ values ];
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class class-name {
              loss-priority (high | low);
            }
          }
        }
      }
      policer cos-policer-name {
        aggregate {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
        premium {
          bandwidth-limit bps;
          burst-size-limit bytes;
        }
      }
    }
  }
```

}

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure Gigabit Ethernet specific interface properties.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Interfaces Overview on page 35
- *gigether-options (ACX Series)*

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default Updating of the ARP cache is disabled on all Ethernet interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gratuitous ARP on page 48
- **no-gratuitous-arp-request on page 609**

 Copyright © 2013, Juniper Networks, Inc.

hold-multiplier

Syntax hold-multiplier *number*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX series switches only) Configure a value for the LLDP hold multiplier.

Hold timer interval in seconds to cache learned LLDP information before discarding.

Options *number*—Advertisement interval multiplier for LLDP cache discard.

Default: 4 (giving 120 second LLDP cache lifetime with other defaults)

Range: 2 through 10

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226

ieee802.1p

Syntax ieee802.1p premium [*values*];

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]

Release Information Statement introduced before Junos Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, configure premium priority values for IEEE 802.1p input traffic.

Options *values*—Define IEEE 802.1p priority values to be treated as premium.

Range: 0 through 7

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Input Priority Map on page 398

ignore-l3-incompletes

Syntax ignore-l3-incompletes;

Hierarchy Level [edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Ignore the counting of Layer 3 incomplete errors on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ignoring Layer 3 Incomplete Errors on page 46

ingress-rate-limit

Syntax ingress-rate-limit *rate*;

Hierarchy Level [edit interfaces *interface-name* fastether-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Perform port-based rate limiting on ingress traffic arriving on Fast Ethernet 8-port, 12-port, and 48-port PICs.

Options *rate*—Traffic rate, in megabits per second (Mbps).

Range: 1 through 100 Mbps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Ingress Rate Limit on page 50

 Copyright © 2013, Juniper Networks, Inc.

inner-tag-protocol-id

Syntax inner-tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.

All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.

Default If the **inner-tag-protocol-id** statement is not configured, the TPID value is 0x8100.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Inner and Outer TPIDs and VLAN IDs on page 203

inner-vlan-id

Syntax inner-vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers or 100-Gigabit Ethernet Type 5 PIC with CFP, or on Ethernet interfaces on EX Series switches, specify the VLAN ID to rewrite for the inner tag of the final packet.

You cannot include the **inner-vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement and the **inner-vlan-id** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **inner-vlan-id** statement you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Options *number*—VLAN ID number.

Range: 0 through 4094

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Inner and Outer TPIDs and VLAN IDs on page 203

 Copyright © 2013, Juniper Networks, Inc.

inner-vlan-id-range

Syntax inner-vlan-id-range start *start-id* end *end-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*],

Release Information Statement introduced in Junos OS Release 9.0.

Description The range of VLAN IDs to be used in the ATM-to-Ethernet interworking cross-connect. Specify the starting VLAN ID and ending VLAN ID.

Options ***start-id***—The lowest VLAN ID to be used.

end-id—The highest VLAN ID to be used.

Range: 32 through 4094

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring ATM-to-Ethernet Interworking*

input-priority-map

Syntax

```
input-priority-map {
  ieee802.1p premium [ values ];
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the input policer priority map to be applied to incoming frames on this interface.

The statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Input Priority Map on page 398
- output-priority-map on page 614

input-vlan-map

See the following sections:

- input-vlan-map (Aggregated Ethernet) on page 584
- input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, 10-Gigabit Ethernet SFP, and 100-Gigabit Ethernet Type 5 PIC with CFP) on page 585

input-vlan-map (Aggregated Ethernet)

Syntax

```
input-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 8.2.

Description

For aggregated Ethernet interfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP only, define the rewrite profile to be applied to incoming frames on this logical interface.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Stacking a VLAN Tag on page 206
- *output-vlan-map (Aggregated Ethernet)*

 Copyright © 2013, Juniper Networks, Inc.

input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, 10-Gigabit Ethernet SFP, and 100-Gigabit Ethernet Type 5 PIC with CFP)

Syntax

```
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;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
pop-pop, **pop-swap**, **push-push**, **swap-push**, and **swap-swap** statements introduced in Junos OS Release 8.1.

Description

For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 10-Gigabit Ethernet SFP interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP only, define the rewrite profile to be applied to incoming frames on this logical interface.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Stacking a VLAN Tag on page 206
- *output-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet with SFPP, 10-Gigabit Ethernet SFP, and 100-Gigabit Ethernet Type 5 PIC with CFP)*

interface

Syntax

```
interface (all | interface-name) {
  disable;
}
```

Hierarchy Level

[edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

(MX Series and T Series routers and EX Series switches only) Specify an LLDP interface.

Options

interface-name—A valid physical interface name.

NOTE: On MX Series and T Series routers, you run LLDP on a physical interface, such as ge-1/0/0, and not at the logical interface (unit) level.

For information about interface names, see *Interface Naming Overview*. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering on TX Matrix routers, see *Routing Matrix with a TX Matrix Router FPC Numbering*.

all—Run LLDP on all interfaces.

disable—Disable LLDP on the specified interface

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226

Copyright © 2013, Juniper Networks, Inc.

interfaces

Syntax interfaces { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure interfaces on the router or switch.

Default The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Physical Interface Configuration Statements Overview*
- Configuring Aggregated Ethernet Link Protection on page 174

lacp

See the following sections:

- lacp (802.3ad) on page 588
- lacp (Aggregated Ethernet) on page 589

lacp (802.3ad)

Syntax

```
lacp {
  traceoptions {
    file lacpd;
    flag all;
  }
  ppm (centralized | distributed);
}
```

Hierarchy Level

[edit interfaces *interface-name* fastether-options 802.3ad],
[edit interfaces *interface-name* gigether-options 802.3ad]

Release Information

Statement introduced in Junos OS Release 9.3.
The **ppm (centralized | distributed)** option introduced in Junos OS Release 9.4.

Description

For aggregated Ethernet interfaces only, configure the Link Aggregation Control Protocol (LACP).

On MX and T Series routers you can specify distributed or centralized periodic packet management (PPM).

Default

If you do not specify **lacp** as either **active** or **passive**, LACP remains passive.

If you do not specify **ppm** as either **centralized** or **distributed**, PPM is distributed.

Options

- **active**—Initiate transmission of LACP packets.
- **passive**—Respond to LACP packets.
- **ppm**—Set PPM to centralized or distributed.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet LACP on page 177

 Copyright © 2013, Juniper Networks, Inc.

lacp (Aggregated Ethernet)

Syntax

```
lacp {
  (active | passive);
  admin-key key;
  accept-data;
  fast-failover;
  link-protection {
    disable;
    (revertive |non-revertive);
  }
  periodic interval;
  system-id mac-address;
  system-priority priority;
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
fast-failover option introduced in Junos OS Release 12.2.

Description For aggregated Ethernet interfaces only, configure the Link Aggregation Control Protocol (LACP).

When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you use the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

Default If you do not specify LACP as either **active** or **passive**, LACP remains passive.

Options **active**—Initiate transmission of LACP packets.

admin-key *number*—Specify an administrative key for the router or switch.

NOTE: You must also configure Multichassis Link Aggregation (MC-LAG) when you configure the **admin-key**.

fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive—Respond to LACP packets.

The remaining statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet LACP on page 177
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

link-discovery

Syntax
link-discovery (active | passive);

Hierarchy Level
[edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 8.2.

Description
For Ethernet interfaces on EX Series switches, and M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

Options
(active | passive)—Passive or active mode. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. Once the discovery process is initiated, both sides participate in discovery.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Link Discovery on page 374

Copyright © 2013, Juniper Networks, Inc.

link-fault-management

Syntax

```
link-fault-management {
  action-profile profile-name {
    action {
      link-down;
      send-critical-event;
      syslog;
    }
    event {
      link-adjacency-loss;
      link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      protocol-down;
    }
  }
  interface interface-name {
    apply-action-profile profile-name;
    link-discovery (active | passive);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    negotiation-options {
      allow-remote-loopback;
      no-allow-link-events;
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers and EX Series switches, specify fault signaling and detection for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 374

Copyright © 2013, Juniper Networks, Inc.

link-mode

Syntax link-mode *mode*;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* ether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Set the device's link connection characteristic.

Options *mode*—Link characteristics:

- **automatic**—Link mode is negotiated. This is the default for EX Series switches.
- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.

Default: Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode. The router's management Ethernet interface, **fxp0** or **em0**, the built-in Fast Ethernet interfaces on the FIC (M7i router), and the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.

NOTE: On J Series ePIM Fast Ethernet interfaces, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."

NOTE: On EX Series switches, if **no-auto-negotiation** is specified in **[edit interfaces *interface-name* ether-options]**, you can select only **full-duplex** or **half-duplex**. If **auto-negotiation** is specified, you can select any mode.

NOTE: Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.

Copyright © 2013, Juniper Networks, Inc.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Link Characteristics on Ethernet Interfaces on page 47
- *Understanding Management Ethernet Interfaces*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

 Copyright © 2013, Juniper Networks, Inc.

link-protection

Syntax

```
link–protection {
  disable;
  (revertive |non-revertive);
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit interfaces ae*x* aggregated-ether-options *lacp*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the **[edit interfaces ge-*fpc/pic/port* gigether-options 802.3ad ae*x*]** hierarchy level or the **[edit interfaces fe-*fpc/pic/port* fastether-options 802.3ad ae*x*]** hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify **link-protection** at the following hierarchy levels:

- **[edit interfaces ge-*fpc/pic/port* ether-options 802.3ad ae*x*]**
- **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]**

For LACP standards link protection, specify **link-protection** at the following hierarchy levels:

- For global LACP link protection, specify at **[edit chassis aggregated-devices ethernet lacp]**
- For a specific aggregated Ethernet interface, specify at **[edit interfaces aeX aggregated-ether-options lacp]**

Options

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 174
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)*

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description For aggregated Ethernet interfaces only, set the required link speed.

Options *speed*—For aggregated Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.

 Copyright © 2013, Juniper Networks, Inc.

- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Aggregated Ethernet Interfaces Overview on page 82
- Configuring Aggregated Ethernet Link Speed on page 186
- Configuring Mixed Aggregated Ethernet Links on page 91
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

Copyright © 2013, Juniper Networks, Inc.

lldp

Syntax

```
lldp {
  advertisement-interval seconds;
  disable;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  port-id-subtype {
    interface-name;
    locally-assigned;
  }
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
}
```

Hierarchy Level [edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX Series switches only) Specify LLDP configuration parameters.

Options The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226

 Copyright © 2013, Juniper Networks, Inc.

lldp-configuration-notification-interval

Syntax lldp-configuration-notification-interval *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX Series switches only) Configure a time for the period of SNMP trap notifications to the Master Agent to wait regarding changes in database information.

Options ***seconds***—Time for the period of SNMP trap notifications about the LLDP database. This feature is disabled by default.

Range: 0 through 3600

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226

Copyright © 2013, Juniper Networks, Inc.

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax (loopback | no-loopback);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces interface-range *name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

NOTE: **By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.**

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet Loopback Capability on page 45

loss-priority

Syntax loss-priority (high | low);

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium forwarding-class *class-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the packet loss priority value.

Options **high**—Packet has high loss priority.

low—Packet has low loss priority.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Output Priority Map on page 399

 Copyright © 2013, Juniper Networks, Inc.

mac-learn-enable

Syntax (mac-learn-enable | no-mac-learn-enable);

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP configure whether source and destination MAC addresses are dynamically learned:

- **mac-learn-enable**—Allow the interface to dynamically learn source and destination MAC addresses.
- **no-mac-learn-enable**—Prohibit the interface from dynamically learning source and destination MAC addresses.

MAC address learning is based on source addresses. You can start accounting for traffic after there has been traffic sent from the MAC address. Once the MAC address is learned, the frames and bytes transmitted to or received from the MAC address can be tracked.

NOTE: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*]** hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see *MX Series Layer 2 Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring MAC Address Filtering on page 401

max-sessions (PPPoE Service Name Tables)

Syntax max-sessions *number*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the maximum number of active PPPoE sessions using either static or dynamic PPPoE interfaces that the router can establish with the specified named service, **empty** service, or **any** service entry in a PPPoE service name table. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

The router uses the **max-sessions** value for a PPPoE service name table entry in conjunction with the **max-sessions** value configured for the PPPoE underlying interface, and with the maximum number of PPPoE sessions supported on your router. If your configuration exceeds any of these maximum session limits, the router is unable to establish the PPPoE session.

Options *number*—Maximum number of active PPPoE sessions that the router can establish with the specified PPPoE service name table entry, in the range 1 to the platform-specific maximum PPPoE sessions supported for your router. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform.

For information about scaling values for PPPoE interfaces, access the *Subscriber Management Scaling Values (XLS)* spreadsheet from the Downloads box on the *Junos OS Subscriber Management* pathway page for the current release.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name on page 508
- Configuring PPPoE Service Name Tables on page 502
- *PPPoE Maximum Session Limit Overview*
- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

max-sessions-vsa-ignore (Static and Dynamic Subscribers)

Syntax max-sessions-vsa-ignore;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the router to ignore (clear) the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143], and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement. The PPPoE maximum session value specifies the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

Default If you do not include the **max-sessions-vsa-ignore** statement, the maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface*
- *PPPoE Maximum Session Limit Overview*
- *Guidelines for Using PPPoE Maximum Session Limit from RADIUS*
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

Copyright © 2013, Juniper Networks, Inc.

maximum-links

Syntax maximum-links *maximum-links-limit*;

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced in Junos OS Release 11.1 for T Series routers.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for MX Series routers.

Description Configure the maximum links limit for aggregated devices.

Options ***maximum-links-limit***—Maximum links limit for aggregated devices.

Range: 16, 32; (PTX Series routers only in Junos OS Release 12.3) 64

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS for Supporting Aggregated Devices on page 88
- Configuring an Aggregated Ethernet Interface on page 87

 Copyright © 2013, Juniper Networks, Inc.

mep

Syntax

```
mep mep-id {
  auto-discovery;
  direction (up | down);
  interface interface-name (protect | working);
  priority number;
  remote-mep mep-id {
    action-profile profile-name;
    sla-iterator-profile profile-name {
      data-tlv-size size;
      iteration-count count-value;
      priority priority-value;
    }
  }
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]

Release Information

Statement introduced in Junos OS Release 8.4.

Description

The numeric identifier of the maintenance association end point (MEP) within the maintenance association.

Options

mep-id—Specify the numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Maintenance Endpoint on page 259
- *Example: Configuring Connectivity Fault Management for a PBB Network on MX Series Routers*

minimum-links

Syntax minimum-links *number*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces ae*x* aggregated-sonet-options],
[edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *range* aggregated-ether-options],
[edit interfaces interface-range *range* aggregated-sonet-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

Options *number*—Number of links.

Range: On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled upOn EX8200 switches, the range of valid values for minimum-links number is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet Minimum Links on page 188
- *Configuring Aggregated SONET/SDH Minimum Links*
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

mip-half-function

Syntax mip-half-function (none | default | explicit);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-managementmaintenance-domain *md-name*],
[edit protocols oam ethernet connectivity-fault-managementmaintenance-association *ma-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Specify the OAM Ethernet CFM maintenance domain MIP half functions.

NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations are the same.

Options **none**—Specify to not use the mip-half-function.

default—Specify to use the default mip-half-function.

explicit—Specify an explicit mip-half-function.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Creating the Maintenance Domain on page 252
- *Example: Configuring Connectivity Fault Management for a PBB Network on MX Series Routers*
- *maintenance-domain*

mpls (Interfaces)

Syntax

```
mpls {
  pop-all-labels {
    required-depth number;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* atm-options],
[edit interfaces *interface-name* sonet-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For passive monitoring on ATM and SONET/SDH interfaces and 10-Gigabit Ethernet interfaces in WAN PHY mode, process incoming IP packets that have MPLS labels.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Removing MPLS Labels from Incoming Packets*
- *Removing MPLS Labels from Incoming Packet*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2013, Juniper Networks, Inc.

no-auto-mdix

Syntax no-auto-mdix;

Hierarchy Level [edit interface ge-*fpc/port/pic* gigether-options]

Release Information Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Disable the Auto MDI/MDIX feature.

MX Series routers with Gigabit Ethernet interfaces automatically detect MDI and MDIX port connections. Use this statement to override the default setting. Remove this statement to return to the default setting.

Default Auto MDI/MDIX is enabled by default.

Options There are no options for this statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Interfaces Overview on page 35
- **gigether-options on page 577**.

no-gratuitous-arp-request

Syntax no-gratuitous-arp-request;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Ethernet interfaces and pseudowire logical interfaces, do not respond to gratuitous ARP requests.

Default Gratuitous ARP responses are enabled on all Ethernet interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gratuitous ARP on page 48
- **gratuitous-arp-reply on page 578**

no-send-pads-ac-info

Syntax no-send-pads-ac-info;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 12.2.

Description Prevent the router from sending the AC-Name and AC-Cookie tags in the PPPoE Active Discovery Session (PADS) packet. When you configure this statement, it affects PADS packets sent on all PPPoE interfaces configured on the router after the command is issued; it has no effect on previously created PPPoE interfaces. By default, the AC-Name and AC-Cookie tags are transmitted in the PADS packet, along with the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags.

NOTE: In Junos OS Release 12.1 and earlier, only the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags are contained in the PADS packet by default. The AC-Name and AC-Cookie tags are not transmitted in the PADS packet by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Disabling the Sending of PPPoE Access Concentrator Tags in PADS Packets on page 511

no-send-pads-error

Syntax no-send-pads-error;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 12.3.

Description Discard PADR messages to prevent transmission of PADS control packets with AC-System-Error tags.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Discarding PADR Messages to Accommodate Abnormal CPE Behavior on page 511

 Copyright © 2013, Juniper Networks, Inc.

oam

Syntax

```
oam {
  ethernet {
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      performance-monitoring {
        delegate-server-processing;
        hardware-assisted-timestamping;
        sla-iterator-profiles {
          profile-name {
            disable;
            calculation-weight {
              delay delay-weight;
              delay-variation delay-variation-weight;
            }
            cycle-time milliseconds;
            iteration-period connections;
            measurement-type (loss | statistical-frame-loss | two-way-delay);
          }
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        name-format (character-string | none | dns | mac+2octet);
        maintenance-association ma-name {
          short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
          protect-maintenance-association protect-ma-name;
          remote-maintenance-association remote-ma-name;
          continuity-check {
            convey-loss-threshold;
            hold-interval minutes;
            interface-status-tlv;
            interval (10m | 10s | 1m | 1s| 100ms);
            loss-threshold number;
            port-status-tlv;
          }
          mep mep-id {
            auto-discovery;
            direction (up | down);
            interface interface-name (protect | working);
            lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
              rem-err-xcon | xcon );
            priority number;
            remote-mep mep-id {
              action-profile profile-name;
```

```
                sla-iterator-profile profile-name {
                  data-tlv-size size;
                  iteration-count count-value;
                  priority priority-value;
                }
              }
            }
          }
        }
      }
      link-fault-management {
        action-profile profile-name {
          action {
            link-down;
            send-critical-event;
            syslog;
          }
          event {
            link-adjacency-loss;
            link-event-rate {
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            protocol-down;
          }
        }
        interface interface-name {
          apply-action-profile
          link-discovery (active | passive);
          pdu-interval interval;
          pdu-threshold threshold-value;
          remote-loopback;
          event-thresholds {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
        }
      }
    }
  }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Copyright © 2013, Juniper Networks, Inc.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers and PTX Series Packet Transport Routers, provide IEEE 802.3ah Operation, Administration, and Maitenance (OAM) support.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 371
- Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers on page 479

optics-options

Syntax

```
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  tx-power dbm;
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
alarm option and **warning** options introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement and **tx-power** option introduced in Junos OS Release 13.2 for PTX Series routers.

Description For 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet DWDM Interface Wavelength Overview on page 441
- *100-Gigabit Ethernet OTN Options Configuration Overview*

Copyright © 2013, Juniper Networks, Inc.

output-priority-map

Syntax

```
output-priority-map {
  classifier {
    premium {
      forwarding-class class-name {
        loss-priority (high | low);
      }
    }
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]
[edit interfaces *interface-name* ether-options ethernet-switch-profile ethernet-policer-profile]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the output policer priority map to be applied to outgoing frames on this interface.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Output Priority Map on page 399
- **input-priority-map on page 583**

 Copyright © 2013, Juniper Networks, Inc.

pado-advertise

Syntax pado-advertise;

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 10.2.

Description Enable named services configured in PPPoE service name tables to be advertised in PPPoE Active Discovery Offer (PADO) control packets. By default, advertisement of named services in PADO packets is disabled.

NOTE: **If you enable advertisement of named services in PADO packets, make sure the number and length of of all advertised service entries does not exceed the maximum transmission unit (MTU) size of the PPPoE underlying interface.**

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Enabling Advertisement of Named Services in PADO Control Packets on page 510

pdu-interval

Syntax pdu-interval *interval*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options *interval*—Periodic OAM PDU sending interval.

Range: 100 through 1000 milliseconds

Default: 1000 milliseconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the OAM PDU Interval on page 375

Copyright © 2013, Juniper Networks, Inc.

pdu-threshold

Syntax pdu-threshold *threshold-value*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options ***threshold-value***—The number of PDUs missed before declaring the peer lost.

Range: 3 through 10 PDUs

Default: 3 PDUs

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the OAM PDU Threshold on page 376

 Copyright © 2013, Juniper Networks, Inc.

periodic

Syntax
periodic *interval*;

Hierarchy Level
[edit interfaces ae*x* aggregated-ether-options lacp],
[edit interfaces interface-range *name* aggregated-ether-options lacp]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description
For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

Options
interval—Interval for periodic transmission of LACP packets.

- **fast**—Transmit packets every second.
- **slow**—Transmit packets every 30 seconds.

Default: fast

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Aggregated Ethernet LACP on page 177
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

Copyright © 2013, Juniper Networks, Inc.

policer

See the following sections:

policer (CFM Firewall)

Syntax

```
policer cfm-policer {
  if-exceeding {
    bandwidth-limit 8k;
    burst-size-limit 2k;
  }
  then discard;
}
```

Hierarchy Level [edit firewall]

Release Information Statement introduced in Junos OS Release 10.0.

Description Attach an explicit policer to CFM sessions.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- policer (CFM Global) on page 619
- policer (CFM Session) on page 620

 Copyright © 2013, Juniper Networks, Inc.

policer (CFM Global)

Syntax

```
policer {
  all cfm-policer-name;
  continuity-check cfm-policer-name;
  other cfm-policer-name;
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in Junos OS Release 10.0.

Description Specify a policer at the global level to police the CFM traffic belonging to all sessions.

Options

continuity-check *cfm-policer-name*—Police all continuity check packets with the policer specified.

other *cfm-policer-name*—Police all non-continuity check packets with the policer specified.

all *cfm-policer-name*—Police all CFM packets with policer specified. If the **all** option is used, then you cannot specify above two options.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- policer (CFM Session) on page 620

policer (CFM Session)

Syntax

```
policer {
  all cfm-policer-name;
  continuity-check cfm-policer-name;
  other cfm-policer-name;
  }
}
```

Hierarchy Level

[edit protocols oam ethernet connectivity-fault-management maintenance-domain *name* level *number* maintenance-association *name*]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Specify a separate policer to rate-limit packets specific to that session.

Options

- **continuity-check *cfm-policer-name***—Police continuity check packets belonging to this session.

  other *cfm-policer-name*—Police all non-continuity check packets belonging to this session.

  all *cfm-policer-name*—Police all CFM packets belonging to this session. If the **all** option is used, then you cannot specify the above two options.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Rate Limiting of Ethernet OAM Messages on page 293
- **policer (CFM Global) on page 619**

Copyright © 2013, Juniper Networks, Inc.

policer (CoS)

Syntax

```
policer cos-policer-name {
  aggregate {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
  premium {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ , Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, define a CoS policer template to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits. The premium policer is not supported on MX Series routers or for Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Options *cos-policer-name*—Name of one policer to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397

policer (MAC)

Syntax

```
policer {
  input cos-policer-name;
  output cos-policer-name;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* accept-source-mac mac-address *mac-address*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* accept-source-mac mac-address *mac-address*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure MAC policing.

NOTE:
On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

Options

input *cos-policer-name*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

output *cos-policer-name*—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring MAC Address Filtering on page 401

Copyright © 2013, Juniper Networks, Inc.

pop

Syntax pop;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, 10-Gigabit Ethernet SFP interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Removing a VLAN Tag on page 207

pop-pop

Syntax pop-pop;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP, and for 10-Gigabit Ethernet SFP interfaces on EX Series switches, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Removing the Outer and Inner VLAN Tags on page 207

Copyright © 2013, Juniper Networks, Inc.

pop-swap

Syntax pop-swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2, IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag on page 208

 Copyright © 2013, Juniper Networks, Inc.

port-id-subtype

Syntax

```
port-id-subtype {
  interface-name;
  locally-assigned;
}
```

Hierarchy Level

[edit protocols lldp]
[edit routing- instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 12.3R1

Description

(MX Series , T Series, and PTX routers only) For Link Layer Discovery Protocol, configure the port ID type, length, and value (TLV).

Options

interface-name—Use the interface name to generate the port ID TLV.

Default: Use the SNMP index of the interface to generate the port ID TLV. This is the default option used to generate port ID TLV.

NOTE: The **show lldp neighbors** command displays the content of the **port ID** TLV received from the peer in the **port Info** field. Changing the configuration of **port-id-subtype** affects the display of the **show lldp neighbors** command on the peer device running Junos OS.

When the value of **port-id-subtype** is set to **locally-assigned**, which is the default value, the **show lldp neighbors** command on the peer device running Junos OS displays the SNMP index as the port information for the local device.

When the value of **port-id-subtype** is set to **interface-name**, the **show lldp neighbors** command on the peer device running Junos OS displays the interface name as the port information for the local device.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **lldp on page 598**
- Configuring LLDP on page 226
- *show lldp neighbors*

port-status-tlv

Syntax port-status-tlv blocked;

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management action-profile *tlv-action* event]

Release Information Statement introduced in Junos OS Release 9.6.

Description Define an **action-profile** consisting of various events and the action. Based on values of **port-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using *action-profile* options.

Options **blocked**—When the incoming CCM packet contains port status TLV with value blocked, the action will be triggered for this action-profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Connectivity Fault Management Action Profile on page 264
- Configuring Remote MEP Action Profile Support on page 285

 Copyright © 2013, Juniper Networks, Inc.

ppp-options

Syntax

```
ppp-options {
  authentication [ authentication-protocols ];
  chap {
    access-profile name;
    challenge-length minimum minimum-length maximum maximum-length;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  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;
  on-demand-ip-address
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
```

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description On interfaces with PPP encapsulation, configure PPP-specific interface properties.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

BEST PRACTICE: On inline service (si) interfaces for L2TP, only the **chap** and **pap** statements are typically used for subscriber management. We recommend that you leave the other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—at their default values.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the PPP Challenge Handshake Authentication Protocol*
- *Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface*

pppoe-options

Syntax

```
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
```

Hierarchy Level

[edit interfaces pp0 unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 7.4.
client Statement introduced in Junos OS Release 8.5.
server Statement introduced in Junos OS Release 8.5.

Description

For J Series Services Routers, M120 Multiservice Edge Routers, M320 Multiservice Edge Service Routers, and MX Series Universal Edge Routers with PPP over Ethernet interfaces, configure PPP over Ethernet-specific interface properties.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a PPPoE Interface on page 497

 Copyright © 2013, Juniper Networks, Inc.

pppoe-underlying-options (Static and Dynamic Subscribers)

Syntax

```
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  duplicate-protection;
  max-sessions number;
  max-sessions-vsa-ignore;
  service-name-table table-name;
  short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 10.0.

Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (**ppp-over-ether**) encapsulation.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE on page 494 (for static interfaces)
- *Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*
- Assigning a Service Name Table to a PPPoE Underlying Interface on page 510

premium

See the following sections:

- premium (Hierarchical Policer) on page 630
- premium (Output Priority Map) on page 631
- premium (Policer) on page 631

premium (Hierarchical Policer)

Syntax

```
premium {
  if-exceeding {
    bandwidth-limit bandwidth;
    burst-size-limit burst;
  }
  then {
    discard;
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* firewall hierarchical-policer],
[edit firewall hierarchical-policer]

Release Information

Statement introduced in Junos OS Release 9.5.
Support at the **[edit dynamic-profiles ... hierarchical-policer *name*]** hierarchy level introduced in Junos OS Release 11.4.

Description

On M40e, M120, and M320 edge routers with FPC input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a premium level for a hierarchical policer.

Options

Options are described separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- *Applying Policers*
- *Junos OS Class of Service Library for Routing Devices*
- *Hierarchical Policer Configuration Overview*
- *Hierarchical Policers*
- *aggregate (Hierarchical Policer)*
- *bandwidth-limit (Hierarchical Policer)*
- *burst-size-limit (Hierarchical Policer)*
- *hierarchical-policer*
- *if-exceeding (Hierarchical Policer)*

 Copyright © 2013, Juniper Networks, Inc.

premium (Output Priority Map)

Syntax

```
premium {
  forwarding-class class-name {
    loss-priority (high | low);
  }
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Specifying an Output Priority Map on page 399
- input-priority-map on page 583

premium (Policer)

Syntax

```
premium {
  bandwidth-limit bps;
  burst-size-limit bytes;
}
```

Hierarchy Level

[edit interfaces *interface-name* gigether-options ethernet-switch-profile ethernet-policer-profile policer *cos-policer-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define a policer to apply to nonpremium traffic.

The statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Policers on page 397
- aggregate (Gigabit Ethernet CoS Policer) on page 554
- ieee802.1p on page 579

Copyright © 2013, Juniper Networks, Inc.

protection-group

Syntax

```
protection-group {
  ethernet-ringring-name {
    control-vlan (vlan-id | vlan-name);
    data-channel {
      vlan number
    }
    east-interface {
      control-channel channel-name {
        vlan number;
      }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    west-interface {
      control-channel channel-name {
        vlan number;
      }
    }
  }
 control-vlan (vlan-id | vlan-name);
    east-interface {
      node-id mac-address;
      control-channel channel-name {
      interface-none
      ring-protection-link-end;
    }
  }
    control-channel channel-name {
      vlan number;
    }
  }
  data-channel {
    vlan number
  }
  guard-interval number;
  node-id mac-address;
  restore-interval number;
  ring-protection-link-owner;
    west-interface {
      node-id mac-address;
      control-channel channel-name {
      interface-none
      ring-protection-link-end;
    }
    control-channel channel-name {
      vlan number;
    }
  }
 }
 guard-interval number;
```

 Copyright © 2013, Juniper Networks, Inc.

```
  restore-interval number;
  traceoptions {
    file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure Ethernet ring protection switching.

The statements are explained separately. All statements apply to MX Series routers. EX Series switches do not assign **node-id** and use **control-vlan** instead of **control-channel**.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

protocol-down

Syntax protocol-down;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in Junos OS Release 8.5.

Description Upper layer indication of protocol down event. When the **protocol-down** statement is included, the protocol down event triggers the action specified under the **action** statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring an OAM Action Profile on page 380

ptopo-configuration-maximum-hold-time

Syntax ptopo-configuration-maximum-hold-time *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX Series switches only) Configure a time to maintain dynamic topology entries.

Options *seconds*—Time to maintain interval dynamic topology entries.

Default: 300

Range: 1 through 2147483647

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring LLDP on page 226

ptopo-configuration-trap-interval

Syntax ptopo-configuration-trap-interval *seconds*;

Hierarchy Level [edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX Series switches only) Configure a time for the period of SNMP trap notifications to the Master Agent to wait regarding changes in topology global statistics.

Options *seconds*—Time for the period of SNMP trap notifications about global statistics. This feature is disabled by default.

Range: 0 through 3600

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring LLDP on page 226

 Copyright © 2013, Juniper Networks, Inc.

push

Syntax push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.

You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP.

If you include the **push** statement in the configuration, you must also include the **pop** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Stacking a VLAN Tag on page 206

push-push

Syntax push-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Stacking Two VLAN Tags on page 209

 Copyright © 2013, Juniper Networks, Inc.

remote-mep

Syntax

```
remote-mep mep-id {
  action-profile profile-name;
  sla-iterator-profile profile-name {
    data-tlv-size size;
    iteration-count count-value;
    priority priority-value;
  }
}
```

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

Options ***mep-id***—Numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately.

Required Privilege Level Configure—To enter configuration mode.
Control—To modify any configuration.

Related Documentation

- Configuring a Maintenance Endpoint on page 259

request

Syntax request (protect | working);

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Perform a manual switch between the protect and working circuits. This statement is honored only if there are no higher-priority reasons to switch.

Options **protect**—Request that the circuit become the protect circuit.

working—Request that the circuit become the working circuit.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Switching Between the Working and Protect Circuits*
- *force*

ring-protection-link-end

Syntax ring-protection-link-end;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name* (east-interface | west-interface)]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

ring-protection-link-owner

Syntax ring-protection-link-owner;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Ethernet Ring Protection Switching Overview on page 527

 Copyright © 2013, Juniper Networks, Inc.

routing-instance (PPPoE Service Name Tables)

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.2.

Description Use in conjunction with the **dynamic-profile** statement at the same hierarchy levels to specify the routing instance in which to instantiate a dynamic PPPoE interface. You can associate a routing instance with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The routing instance associated with a service entry in a PPPoE service name table overrides the routing instance associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **routing-instance** statement at the **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level, you cannot also include the **static-interface** statement at this level. The **routing-instance** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options *routing-instance-name*—Name of the routing instance in which the router instantiates the dynamic PPPoE interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- *Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

Copyright © 2013, Juniper Networks, Inc.

sa-multicast (100-Gigabit Ethernet)

Syntax sa-multicast;

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the 100-Gigabit Ethernet PIC or MIC to interoperate with other Juniper Networks 100-Gigabit Ethernet PICs.

NOTE: **The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.**

sa-multicast supports interoperability between the following PICs and MICs:

- 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP) and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) .
- 100-Gigabit Ethernet MICs and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP on page 464
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 on page 465
- *Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4)(Type 4 1X100GE PIC for STFPC4 FPC) Using SA Multicast Mode*
- *Interoperability Between MPC4E (MPC4E-3D-2CGE-8XGE) and 100-Gigabit Ethernet PICs on Type 4 FPC*
- *Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode*
- *Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP*
- *Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4*
- **forwarding-mode (100-Gigabit Ethernet) on page 575**
- *sa-multicast (PTX Series Packet Transport Routers)*

 Copyright © 2013, Juniper Networks, Inc.

- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 684**
- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 460

service (PPPoE)

Syntax

```
service service-name {
  drop;
  delay seconds;
  terminate;
  dynamic-profile profile-name;
  routing-instance routing-instance-name;
  max-sessions number;
  agent-specifier {
    aci circuit-id-string ari remote-id-string {
      drop;
      delay seconds;
      terminate;
      dynamic-profile profile-name;
      routing-instance routing-instance-name;
      static-interface interface-name;
    }
  }
}
```

Hierarchy Level [edit protocols pppoe service-name-tables *table-name*]

Release Information Statement introduced in Junos OS Release 10.0.
any, **dynamic-profile**, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Specify the action taken by the interface on receipt of a PPPoE Active Discovery Initiation (PADI) control packet for the specified named service, **empty** service, or **any** service in a PPPoE service name table. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service.

Default The default action is terminate.

Options *service-name*—Service entry in the PPPoE service name table:

- *service-name*—Named service entry of up to 32 characters; for example, **premiumService**. You can configure a maximum of 512 named service entries across all PPPoE service name tables on the router.
- **empty**—Service entry of zero length that represents an unspecified service. Each PPPoE service name table includes one **empty** service entry by default.
- **any**—Default service for non-empty service entries that do not match the named or **empty** service entries configured in the PPPoE service name table. Each PPPoE service name table includes one **any** service entry by default.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag on page 506
- Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag on page 504
- Configuring the Action Taken for the Any Service on page 505

service-name-table

Syntax service-name-table *table-name*;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information Statement introduced in Junos OS Release 10.0.
Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

Description Specify the PPPoE service name table assigned to a PPPoE underlying interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive.

NOTE: The **[edit ... family pppoe]** hierarchies are supported only on MX Series routers with MPCs.

Options *table-name*—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Assigning a Service Name Table to a PPPoE Underlying Interface on page 510
- *Configuring the PPPoE Family for an Underlying Interface*

service-name-tables

Syntax

```
service-name-tables table-name {
  service service-name {
    drop;
    delay seconds;
    terminate;
    dynamic-profile profile-name;
    routing-instance routing-instance-name;
    max-sessions number;
    agent-specifier {
      aci circuit-id-string ari remote-id-string {
        drop;
        delay seconds;
        terminate;
        dynamic-profile profile-name;
        routing-instance routing-instance-name;
        static-interface interface-name;
      }
    }
  }
}
```

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 10.0.
dynamic-profile, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

Description Create and configure a PPPoE service name table. Specify the action taken for each service and remote access concentrator on receipt of a PPPoE Active Discovery Initiation (PADI) packet. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service. A maximum of 32 PPPoE service name tables is supported per router.

Options ***table-name***—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Creating a Service Name Table on page 503

 Copyright © 2013, Juniper Networks, Inc.

short-cycle-protection (Static and Dynamic Subscribers)

Syntax short-cycle-protection <lockout-time-min *minimum-seconds* lockout-time-max *maximum-seconds*>;

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit interfaces demux0 unit *logical-unit-number* family pppoe]
[edit interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family pppoe],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the router to temporarily prevent (lock out) a failed or short-lived (also known as short-cycle) PPPoE subscriber session from reconnecting for a default or configurable period of time. You can optionally override the default lockout time, 1 through 300 seconds (5 minutes), by specifying the minimum lockout time and maximum lockout time as part of the **short-cycle-protection** statement.

Options **lockout-time-min *minimum-seconds***—(Optional) Minimum lockout time for failed or short-lived PPPoE subscriber sessions. The *minimum-seconds* value must be less than or equal to the *maximum-seconds* value. Setting *minimum-seconds* and *maximum-seconds* to the same value causes the lockout time to become fixed at that value.

Range: 1 through 86400 (24 hours)

Default: 1

lockout-time-max *maximum-seconds*—(Optional) Maximum lockout time for failed or short-lived PPPoE subscriber sessions. The *maximum-seconds* value must be equal to or greater than the *minimum-seconds* value. Setting *maximum-seconds* and *minimum-seconds* to the same value causes the lockout time to become fixed at that value.

Range: 1 through 86400 (24 hours)

Default: 300 (5 minutes)

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Lockout of PPPoE Subscriber Sessions*
- *PPPoE Subscriber Session Lockout Overview*
- *Understanding the Lockout Period for PPPoE Subscriber Session Lockout*

- For information about configuring dynamic PPPoE subscriber interfaces, see the *Junos OS Subscriber Management and Services Library*
- For information about configuring static PPPoE interfaces, see the *Ethernet Interfaces*

 Copyright © 2013, Juniper Networks, Inc.

source-address-filter

Syntax

```
source-address-filter {
  mac-address;
}
```

Hierarchy Level

[edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the **source-filtering** statement in the configuration to enable source address filtering. This statement is not supported on the J Series Services Routers.

Options

mac-address—MAC address filter. You can specify the MAC address as ***nn*:*nn*:*nn*:*nn*:*nn*:*nn*** or ***nnnn*.*nnnn*.*nnnn***, where ***n*** is a decimal digit. To specify more than one address, include multiple ***mac-address*** options in the **source-address-filter** statement.

If you enable the VRRP on a Fast Ethernet or Gigabit Ethernet interface, as described in "VRRP and VRRP for IPv6 Overview" on page 391, and if you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 3768, *Virtual Router Redundancy Protocol*. When you configure the VRRP group, the group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Ethernet MAC Address Filtering on page 42
- Configuring MAC Filtering on PTX Series Packet Transport Routers on page 475
- **source-filtering on page 648**

source-filtering

Syntax (source-filtering | no-source-filtering);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the **source-address-filter** statement.

If the remote Ethernet card is changed, the interface is no longer able to receive packets from the new card because it has a different MAC address.

Default Source address filtering is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Enabling Ethernet MAC Address Filtering on page 42
- Configuring MAC Filtering on PTX Series Packet Transport Routers on page 475
- *accept-source-mac*
- **source-address-filter on page 647**

 Copyright © 2013, Juniper Networks, Inc.

speed

See the following sections:

- speed (Ethernet) on page 650
- speed (MX Series DPC) on page 651

Copyright © 2013, Juniper Networks, Inc.

speed (Ethernet)

Syntax speed (10m | 100m | 1g | auto | auto-10m-100m);

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure the interface speed. This statement applies to the management Ethernet interface (**fxp0** or **em0**), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), the built-in Ethernet interfaces on J Series Services Routers, Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode, and Gigabit Ethernet interfaces on EX Series switches.

When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.

Options You can specify the speed as either **10m** (10 Mbps), **100m** (100 Mbps), or on J Series routers with uPIMs installed and on MX Series routers, **1g** (1 Gbps). You can also specify the **auto** option on MX Series routers.

For Gigabit Ethernet interfaces on EX Series switches, you can specify one of the following options:

- **10m**—10 Mbps
- **100m**—100 Mbps
- **1g**—1 Gbps
- **auto**—Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link.
- **auto-10m-100m**—Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Interface Speed*
- Configuring the Interface Speed on Ethernet Interfaces on page 49
- Configuring Gigabit Ethernet Autonegotiation on page 416
- Configuring J Series Services Router Switching Interfaces on page 40
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

Copyright © 2013, Juniper Networks, Inc.

speed (MX Series DPC)

Syntax speed (auto | 1Gbps | 100Mbps | 10Mbps);

Hierarchy Level [edit interfaces ge-*fpc/pic/port*]

Release Information Statement introduced in Junos OS Release 9.5.

Description On MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs you can set auto negotiation of speed. To specify the auto negotiation speed, use the **speed (auto | 1Gbps | 100Mbps | 10Mbps)** statement under the **[edit interface ge-/fpc/pic/port]** hierarchy level. The **auto** option will attempt to automatically match the rate of the connected interface. To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**.

NOTE: If the negotiated speed and the interface speed do not match, the link will not be brought up. Half duplex mode is not supported.

Options You can specify the speed as either **auto** (autonegotiate), **10Mbps** (10 Mbps), **100Mbps** (100 Mbps), or **1Gbps** (1 Gbps).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Autonegotiation on page 416
- no-auto-mdix on page 609

static-interface

Syntax

static-interface *interface-name*;

Hierarchy Level

[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Reserve the specified static PPPoE interface for use only by the PPPoE client with matching agent circuit identifier (ACI) and agent remote identifier (ARI) information. You can specify only one static interface per ACI/ARI pair configured for a named service entry, **empty** service entry, or **any** service entry in the PPPoE service name table.

The static interface associated with an ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.

If you include the **static-interface** statement in the configuration, you cannot also include either the **dynamic-profile** statement or the **routing-instance** statement. The **dynamic-profile**, **routing-instance**, and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

Options

interface-name—Name of the static PPPoE interface reserved for use by the PPPoE client with matching ACI/ARI information. Specify the interface in the format **pp0.*logical***, where ***logical*** is a logical unit number from 0 through 16385 for static interfaces.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client on page 509

 Copyright © 2013, Juniper Networks, Inc.

swap

Syntax swap;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.

On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rewriting the VLAN Tag on Tagged Frames on page 210

swap-push

Syntax swap-push;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag on page 214

 Copyright © 2013, Juniper Networks, Inc.

swap-swap

Syntax

swap-swap;

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.

You can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and for 100-Gigabit Ethernet Type 5 PIC with CFP.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rewriting the Inner and Outer VLAN Tags on page 214

switch-options

Syntax

```
switch-options {
   switch-port port-number {
      (auto-negotiation | no-auto-negotiation);
      speed (10m | 100m | 1g);
      link-mode (full-duplex | half-duplex);
   }
}
```

Hierarchy Level [edit interfaces ge-*pim*/0/0]

Release Information Statement introduced in Junos OS Release 8.4.

Description On a J Series Services Router with multiport Gigabit Ethernet uPIMs installed and operating in access switching mode, only one physical interface is configured for the entire multiport Gigabit Ethernet uPIM. Configuration of the physical port characteristics is done under the single physical interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring J Series Services Router Switching Interfaces on page 40

Copyright © 2013, Juniper Networks, Inc.

switch-port

Syntax

```
switch-port port-number {
  (auto-negotiation | no-auto-negotiation);
  speed (10m | 100m | 1g);
  link-mode (full-duplex | half-duplex);
}
```

Hierarchy Level [edit interfaces ge-*pim*/0/0 switch-options]

Release Information Statement introduced in Junos OS Release 8.4.

Description On a J Series Services Router with Ethernet uPIMs installed and operating in access switching mode, configuration of the physical port characteristics, done under the single physical interface.

Default Autonegotiation is enabled by default. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation.

Options ***port-number***—Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring J Series Services Router Switching Interfaces on page 40

system-id

Syntax system-id *system-id*;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp]

Release Information Statement introduced in Junos OS Release 12.2R1

Description Define the LACP system identifier at the aggregated Ethernet interface level.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet LACP on page 177

 Copyright © 2013, Juniper Networks, Inc.

tag-protocol-id

See the following sections:

- tag-protocol-id (TPIDs Expected to Be Sent or Received) on page 659
- tag-protocol-id (TPID to Rewrite) on page 660

tag-protocol-id (TPIDs Expected to Be Sent or Received)

Syntax tag-protocol-id [*tpids*];

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile],
[edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the **tag-protocol-id** statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (**0x8100**) is supported.

Options *tpids*—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames on page 201

tag-protocol-id (TPID to Rewrite)

Syntax tag-protocol-id *tpid*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the **[edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]]** hierarchy level.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers the default TPID value (**0x8100**) is supported.

Default If the **tag-protocol-id** statement is not configured, the TPID value is 0x8100.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Inner and Outer TPIDs and VLAN IDs on page 203

 Copyright © 2013, Juniper Networks, Inc.

terminate (PPPoE Service Name Tables)

Syntax terminate;

Hierarchy Level [edit protocols pppoe service-name-tables *table-name* service *service-name*],
[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]

Release Information Statement introduced in Junos OS Release 10.0.
Support at **[edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*]** hierarchy level introduced in Junos OS Release 10.2.

Description Direct the router to immediately respond to a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client by sending the client a PPPoE Active Discovery Offer (PADO) packet. The PADO packet contains the name of the access concentrator (router) that can service the client request. The **terminate** action is the default action for a named service entry, **empty** service entry, **any** service entry, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in a PPPoE service name table.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring PPPoE Service Name Tables on page 502

traceoptions (Protocols LLDP)

Syntax

```
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
}
```

Hierarchy Level

[edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Set LLDP protocol-level tracing options.

Default

The default LLDP protocol-level trace options are inherited from the global **traceoptions** statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place spanning-tree protocol tracing output in the file **/var/log/stp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the LLDP-specific tracing options:

- **all**—Trace all operations.
- **config**—Log configuration events.
- **interface**—Trace interface update events.
- **protocol**—Trace protocol information.
- **rtsock**—Trace socket events.
- **vlan**—Trace vlan update events.

Copyright © 2013, Juniper Networks, Inc.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events. This is the default. If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file. This is the default. If you do not include this option, tracing output is appended to an existing trace file.

size *maximum-file-size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Tracing LLDP Operations on page 228

traceoptions (PPPoE)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
  filter {
    aci regular-expression;
    ari regular-expresion;
    service-name regular-expresion;
    underlying-interface interface-name;
  }
  flag flag;
  level (all | error | info | notice | verbose | warning);
  no-remote-trace;
}
```

Hierarchy Level [edit protocols pppoe]

Release Information Statement introduced in Junos OS Release 9.6.
Option **filter** introduced in Junos OS Release 12.3

Description Define tracing operations for PPPoE processes.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

disable—Disable this trace flag.

filter—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

BEST PRACTICE: **Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.**

- **aci *regular-expression***—Regular expression to match the agent circuit identifier provided by PPPoE client.
- **ari *regular-expression***—Regular expression to match the agent remote identifier provided by PPPoE client.

 Copyright © 2013, Juniper Networks, Inc.

- **service** ***regular-expression***—Regular expression to match the name of PPPoE service.
- **underlying-interface** ***interface-name***—Name of a PPPoE underlying interface. You cannot use a regular expression for this filter option.

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **config**—Trace configuration events.
- **events**—Trace events.
- **gres**—Trace GRES events.
- **init**—Trace initialization events.
- **interface-db**—Trace interface database operations.
- **memory**—Trace memory processing events.
- **protocol**—Trace protocol events.
- **rtsock**—Trace routing socket events.
- **session-db**—Trace connection events and flow.
- **signal**—Trace signal operations.
- **state**—Trace state handling events.
- **timer**—Trace timer processing.
- **ui**—Trace user interface processing.

level—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size ***maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***size*k** to specify KB, ***size*m** to specify MB, or ***size*g** to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace—To view this statement in the configuration.
trace-control—To add this statement to the configuration.

Related Documentation

- Configuring PPPoE Service Name Tables on page 502
- Tracing PPPoE Operations on page 514

transmit-delay

Syntax

transmit-delay *seconds*;

Hierarchy Level

[edit protocols lldp],
[edit routing-instances *routing-instance-name* protocols lldp]

Release Information

Statement introduced in Junos OS Release 9.6.

Description

(MX Series and T Series routers and EX Series switches only) Configure a delay between two successive LLDP advertisements.

Options

seconds—Delay between two successive LLDP advertisements.

Default: 2

Range: 1 through 8192

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP on page 226

 Copyright © 2013, Juniper Networks, Inc.

unit

Syntax

```
unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  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 {
      maximum-contexts number <force>;
      f-max-period number;
      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;
family family-name {
  ... the family subhierarchy appears after the main [edit interfaces interface-name unit
    logical-unit-number] hierarchy ...
}
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 {
  up-count cells;
  down-count cells;
}
oam-period (disable | seconds);
```

 Copyright © 2013, Juniper Networks, Inc.

```
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;
  }
  dynamic-profile profile-name;
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-restart-timer milliseconds;
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  max-sessions number;
}
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;
targeted-distribution;
```

Copyright © 2013, Juniper Networks, Inc.

```
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-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      (input | output | input output);
    }
  }
  access-concentrator name;
  address address {
    ... the address subhierarchy appears after the main [edit interfaces interface-name unit
      logical-unit-number family family-name] hierarchy ...
  }
  bridge-domain-type (bvlan | svlan);
  bundle interface-name;
  core-facing;
  demux-destination {
    destination-prefix;
  }
  demux-source {
    source-prefix;
  }
  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 ];
  }
  interface-mode (access | trunk);
  ipsec-sa sa-name;
  isid-list all-service-groups;
  keep-address-and-control;
  mac-validate (loose | strict);
  max-sessions number;
```

 Copyright © 2013, Juniper Networks, Inc.

```
mtu bytes;
multicast-only;
no-redirects;
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
primary;
protocols [inet iso mpls];
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 {
  input;
  output;
}
service {
  input {
    post-service-filter filter-name;
    service-set service-set-name <service-filter filter-name>;
  }
  output {
    service-set service-set-name <service-filter filter-name>;
  }
}
service-name-table table-name
(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;
vlan-id number;
vlan-id-list [number number-number];
address address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  master-only;
  multipoint-destination address {
    dlci dlci-identifier;
    epd-threshold cells <plp1 cells>;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
```

```
                (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
                    sustained rate);
                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;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost number;
                }
                priority-hold-time seconds;
                route ip-address/prefix-length routing-instance instance-name priority-cost cost;
            }
            virtual-address [ addresses ];
            virtual-link-local-address ipv6–address;
            vrrp-inherit-from {
                active-interface interface-name;
                active-group group-number;
            }
        }
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*],
[edit interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number***—Number of the logical unit.

Range: 0 through 1,073,741,823 for demux and PPPoE static interfaces only. 0 through 16,385 for all other static interface types.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

 Copyright © 2013, Juniper Networks, Inc.

Related Documentation

- *Configuring Logical Interface Properties*
- *Example: Configuring E-LINE and E-LAN Services for a PBB Network on MX Series Routers*
- *Junos OS Services Interfaces Library for Routing Devices*

version-3

Syntax version-3;

Hierarchy Level [edit protocols vrrp]

Release Information Statement introduced in Junos OS Release 12.2.

Description Enable Virtual Router Redundancy Protocol version 3 (VRRPv3).

NOTE:

- Even though the **version-3** statement can be configured only at the **[edit protocols vrrp]** hierarchy level, VRRPv3 is enabled on all the configured logical systems as well.
- When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Support for VRRPv3*
- *VRRP Configuration Hierarchy*
- *VRRP for IPv6 Configuration Hierarchy*

vlan-id

See the following sections:

vlan-id (Logical Port in Bridge Domain)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in Junos OS Release 9.2.

Description The VLAN ID configured on the logical port. Received packets with no VLAN tags are forwarded within the bridge domain with the matching VLAN ID.

Options **number**—The VLAN ID.

Range: 1 through 4095

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Logical Interface for Access Mode on page 75

 Copyright © 2013, Juniper Networks, Inc.

vlan-id (Outer VLAN ID)

Syntax vlan-id *outer-vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.0.

Description The outer VLAN ID to be used in ATM-to-Ethernet interworking cross-connects. Outer VLAN IDs are converted to the ATM VPI. The outer VLAN ID must match the VPI value configured. The allowable VPI range is 0 to 255. Do not configure the outer VLAN ID to be greater than 255.

Options **outer-vlan-id**—Outer VLAN ID number.

Range: 0 through 4094

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring ATM-to-Ethernet Interworking*

vlan-id (VLAN ID to Be Bound to a Logical Interface)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options *number*—A valid VLAN identifier.

Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.

For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.

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

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Mixed Tagging on page 57

vlan-id (VLAN ID to Rewrite)

Syntax vlan-id *number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.

You cannot include the **vlan-id** statement with the **swap** statement, **swap-push** statement, **push-push** statement, or **push-swap** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]** hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the **vlan-id** statement that you include at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rewriting the VLAN Tag on Tagged Frames on page 210
- Binding VLAN IDs to Logical Interfaces on page 59

 Copyright © 2013, Juniper Networks, Inc.

vlan-id-list

See the following sections:

- vlan-id-list (Ethernet VLAN Circuit) on page 678
- vlan-id-list (Interface in Bridge Domain) on page 679

Copyright © 2013, Juniper Networks, Inc.

vlan-id-list (Ethernet VLAN Circuit)

Syntax vlan-id-list [*vlan-id vlan-id–vlan-id*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description (MX Series routers only) Binds a single-tag logical interface to a list of VLAN IDs. Configures a logical interface to receive and forward any tag frame whose VLAN ID tag matches the list of VLAN IDs you specify.

NOTE:
When you create a circuit cross-connect (CCC) using VLAN-bundled single-tag logical interfaces on Layer 2 VPN routing instances, the circuit automatically uses **ethernet** encapsulation. For Layer 2 VPN, you need to include the **encapsulation-type** statement and specify the value **ethernet** at either of the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols l2vpn]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]**

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *Junos OS VPNs Library for Routing Devices*.

Options **[*vlan-id vlan-id–vlan-id*]**—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring **vlan-id-list** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id-range 1-4094;
}
```

 Copyright © 2013, Juniper Networks, Inc.

```
[edit interfaces interface-name]
unit 0;
```

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Binding VLAN IDs to Logical Interfaces on page 59
- *encapsulation (Logical Interface)*
- *encapsulation (Physical Interface)*
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*
- *flexible-vlan-tagging*
- **vlan-tagging on page 685**
- **vlan-tags (Dual-Tagged Logical Interface) on page 687**

vlan-id-list (Interface in Bridge Domain)

Syntax vlan-id-list [*number number-number*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family bridge],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.

Options ***number number***—Individual VLAN IDs separated by a space.

number-number—Starting VLAN ID and ending VLAN ID in an inclusive range.

Range: 1 through 4095

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Logical Interface for Trunk Mode on page 76
- Configuring the VLAN ID List for a Trunk Interface on page 77

Copyright © 2013, Juniper Networks, Inc.

vlan-id-range

Syntax vlan-id-range *vlan-id–vlan-id*

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Bind a range of VLAN IDs to a logical interface.

Options **number**—The first number is the lowest VLAN ID in the range the second number is the highest VLAN ID in the range.

Range: 1 through 4094

NOTE: Configuring **vlan-id-range** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id-range 1-4094;
}
```

```
[edit interfaces interface-name]
unit 0;
```

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

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Binding a Range of VLAN IDs to a Logical Interface on page 61

 Copyright © 2013, Juniper Networks, Inc.

vlan-ranges

Syntax

```
vlan-ranges {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-82 <circuit-id> <remote-id>;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | inet);
    ranges (any | low-tag)–(any | high-tag);
  }
  override;
}
```

Hierarchy Level [edit interfaces *interface-name* auto-configure]

Release Information Statement introduced in Junos OS Release 9.5.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- *Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles*
- *Configuring Dynamic Mixed VLAN Ranges*

Copyright © 2013, Juniper Networks, Inc.

vlan-rewrite

Syntax vlan-rewrite translate (200 500 | 201 501)

Hierarchy Level [edit interfaces *interface-name* unit *number* family bridge interface-mode trunk]
[edit interfaces *interface-name* unit *number* family ethernet-switching interface-mode trunk]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.

Options **translate 200 500**—Translates incoming packets with VLAN 200 to 500.

translate 201 501—Translates incoming packets with VLAN 201 to 501.

translate 202 502—Translates incoming packets with VLAN 202 to 502.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Rewriting a VLAN Tag and Adding a New Tag on page 214

 Copyright © 2013, Juniper Networks, Inc.

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)

Syntax vlan-rule (high-low | odd-even);

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode vlan-steering]

Release Information Statement introduced in Junos OS Release 10.4.

Description Configure the interoperation mode of the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) when interoperating with 100 gigabit Ethernet interfaces from other vendors.

If no VLAN rule is configured, all tagged packets are distributed to PFE0.

Options **high-low**—VLAN IDs 1 through 2047 are distributed to PFE0 and VLAN IDs 2048 through 4096 are distributed to PFE1.

odd-even—Odd number VLAN IDs are distributed to PFE1 and even number VLAN IDs are distributed to PFE0.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 460
- **forwarding-mode (100-Gigabit Ethernet) on page 575**
- **vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) on page 684**

Copyright © 2013, Juniper Networks, Inc.

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)

Syntax

```
vlan-steering {
  vlan-rule (high-low | odd-even);
}
```

Hierarchy Level [edit chassis fpc *slot* pic *slot* forwarding-mode]

Release Information Statement introduced in Junos OS Release 9.4.

Description Configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with 100 gigabit Ethernet interfaces from other vendors.

The other statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP on page 460
- forwarding-mode (100-Gigabit Ethernet) on page 575
- sa-multicast (100-Gigabit Ethernet) on page 640
- vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) on page 683

 Copyright © 2013, Juniper Networks, Inc.

vlan-tagging

Syntax vlan-tagging;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.

Description For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *Example: Configuring BGP Autodiscovery for LDP VPLS*
- *Configuring a Layer 3 Subinterface (CLI Procedure)*
- *Configuring Tagged Aggregated Ethernet Interfaces*
- *Configuring Interfaces for VPLS Routing*
- *Enabling VLAN Tagging*
- 802.1Q VLANs Overview on page 53
- *vlan-id*

Copyright © 2013, Juniper Networks, Inc.

vlan-tags

See the following sections:

- vlan-tags (Dual-Tagged Logical Interface) on page 687
- vlan-tags (Stacked VLAN Tags) on page 689

 Copyright © 2013, Juniper Networks, Inc.

vlan-tags (Dual-Tagged Logical Interface)

Syntax vlan-tags inner-list [*vlan-id vlan-id–vlan-id*] outer <*tpid*.>*vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.5.

Description (MX Series routers only) Binds a dual-tag logical interface to a list of VLAN IDs. Configures the logical interface to receive and forward any dual-tag frame whose inner VLAN ID tag matches the list of VLAN IDs you specify.

NOTE:

To create a circuit cross-connect (CCC) using VLAN-bundled dual-tag logical interfaces on Layer 2 VPN routing instances, you must include the **encapsulation-type** statement and specify the value **ethernet-vlan** at the one of the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols l2vpn]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]**

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *Junos OS VPNs Library for Routing Devices*.

Options **inner-list [*vlan-id vlan-id vlan-id-vlan-id*]**—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer <*tpid*.>*vlan-id*—An optional Tag Protocol ID (TPID) and a valid VLAN ID.

Range: For TPID, specify a hexadecimal value in the format **0x*nnnn***.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Configuring **inner-list** with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1 through 4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-tags outer vid inner-list 1-4094;
}
```

```
[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Binding VLAN IDs to Logical Interfaces on page 59
- *encapsulation (Logical Interface)*
- *encapsulation (Physical Interface)*
- encapsulation-type (Layer 2 VPN routing instance), see the *Junos OS VPNs Library for Routing Devices*.
- *flexible-vlan-tagging*
- **vlan-id-list (Ethernet VLAN Circuit) on page 678**
- **vlan-tagging on page 685**

 Copyright © 2013, Juniper Networks, Inc.

vlan-tags (Stacked VLAN Tags)

Syntax vlan-tags inner *tpid.vlan-id* inner-range *vid1—vid2* outer *tpid.vlan-id*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Description For Gigabit Ethernet IQ and IQE interfaces only, bind TPIDs and 802.1Q VLAN tag IDs to a logical interface.

NOTE: The **inner-range** *vid1–vid2* option is supported on MX Series with IQE PICs only.

Options **inner** *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

Range: (PTX Series) For VLAN ID, 0 through 4094.

inner-range *vid1–vid2*—For MX Series routers with Enhanced IQ (IQE) PICs only; specify a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: (most routers) For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.

Range: (PTX Series) For VLAN ID, 0 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces.

NOTE: Configuring **inner-range** with the entire vlan-id range consumes system resources and is not a best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1–4094), it has the same result as not specifying a range; however, it consumes Packet Forwarding Engine resources such as VLAN lookup table entries, and so on.

The following examples illustrate this further:

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

```
stacked-vlan-tagging;
unit number {
    vlan-tags outer vid inner-range 1-4094;
}

[edit interfaces interface-name]
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Dual VLAN Tags on page 202
- Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers on page 476
- *stacked-vlan-tagging*

vlan-tags-outer

Syntax vlan-tags-outer *vlan-tag*;

Hierarchy Level [edit interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.5.

Description The S-VLAN outer tag that belongs to a set of interfaces used to configure hierarchical CoS schedulers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Class of Service Library for Routing Devices*

 Copyright © 2013, Juniper Networks, Inc.

vlan-vci-tagging

Syntax vlan-vci-tagging;

Hierarchy Level [edit interfaces *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable the ATM-to-Ethernet interworking cross-connect function on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring ATM-to-Ethernet Interworking*

Copyright © 2013, Juniper Networks, Inc.

wavelength

Syntax wavelength *nm*;

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description For 10-Gigabit or 100-Gigabit Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.

Options *nm*—Wavelength value. It can be one of the following:

- **1528.77**—1528.77 nanometers (nm), corresponds to 50 GHz through 100 GHz
- **1529.16**—1529.16 nm, corresponds to 50 GHz
- **1529.55**—1529.55 nm, corresponds to 50 GHz through 100 GHz
- **1529.94**—1529.94 nm, corresponds to 50 GHz
- **1530.33**—1530.33 nm, corresponds to 50 GHz through 100 GHz
- **1530.72**—1530.72 nm, corresponds to 50 GHz
- **1531.12**—1531.12 nm, corresponds to 50 GHz through 100 GHz
- **1531.51**—1531.51 nm, corresponds to 50 GHz
- **1531.90**—1531.90 nm, corresponds to 50 GHz through 100 GHz
- **1532.29**—1532.29 nm, corresponds to 50 GHz
- **1532.68**—1532.68 nm, corresponds to 50 GHz through 100 GHz
- **1533.07**—1533.07 nm, corresponds to 50 GHz
- **1533.47**—1533.47 nm, corresponds to 50 GHz through 100 GHz
- **1533.86**—1533.86 nm, corresponds to 50 GHz
- **1534.25**—1534.25 nm, corresponds to 50 GHz through 100 GHz
- **1534.64**—1534.64 nm, corresponds to 50 GHz
- **1535.04**—1535.04 nm, corresponds to 50 GHz through 100 GHz
- **1535.43**—1535.43 nm, corresponds to 50 GHz
- **1535.82**—1535.82 nm, corresponds to 50 GHz through 100 GHz
- **1536.22**—1536.22 nm, corresponds to 50 GHz
- **1536.61**—1536.61 nm, corresponds to 50 GHz through 100 GHz
- **1537.00**—1537.00 nm, corresponds to 50 GHz
- **1537.40**—1537.40 nm, corresponds to 50 GHz through 100 GHz

Copyright © 2013, Juniper Networks, Inc.

- **1537.79**—1537.79 nm, corresponds to 50 GHz
- **1538.19**—1538.19 nm, corresponds to 50 GHz through 100 GHz
- **1538.58**—1538.58 nm, corresponds to 50 GHz
- **1538.98**—1538.98 nm, corresponds to 50 GHz through 100 GHz
- **1539.37**—1539.37 nm, corresponds to 50 GHz
- **1539.77**—1539.77 nm, corresponds to 50 GHz through 100 GHz
- **1540.16**—1540.16 nm, corresponds to 50 GHz
- **1540.56**—1540.56 nm, corresponds to 50 GHz through 100 GHz
- **1540.95**—1540.95 nm, corresponds to 50 GHz
- **1541.35**—1541.35 nm, corresponds to 50 GHz through 100 GHz
- **1541.75**—1541.75 nm, corresponds to 50 GHz
- **1542.14**—1542.14 nm, corresponds to 50 GHz through 100 GHz
- **1542.54**—1542.54 nm, corresponds to 50 GHz
- **1542.94**—1542.94 nm, corresponds to 50 GHz through 100 GHz
- **1543.33**—1543.33 nm, corresponds to 50 GHz
- **1543.73**—1543.73 nm, corresponds to 50 GHz through 100 GHz
- **1544.13**—1544.13 nm, corresponds to 50 GHz
- **1544.53**—1544.53 nm, corresponds to 50 GHz through 100 GHz
- **1544.92**—1544.92 nm, corresponds to 50 GHz
- **1545.32**—1545.32 nm, corresponds to 50 GHz through 100 GHz
- **1545.72**—1545.72 nm, corresponds to 50 GHz
- **1546.12**—1546.12 nm, corresponds to 50 GHz through 100 GHz
- **1546.52**—1546.52 nm, corresponds to 50 GHz
- **1546.92**—1546.92 nm, corresponds to 50 GHz through 100 GHz
- **1547.32**—1547.32 nm, corresponds to 50 GHz
- **1547.72**—1547.72 nm, corresponds to 50 GHz through 100 GHz
- **1548.11**—1548.11 nm, corresponds to 50 GHz
- **1548.51**—1548.51 nm, corresponds to 50 GHz through 100 GHz
- **1548.91**—1548.91 nm, corresponds to 50 GHz
- **1549.32**—1549.32 nm, corresponds to 50 GHz through 100 GHz
- **1549.72**—1549.72 nm, corresponds to 50 GHz
- **1550.12**—1550.12 nm, corresponds to 50 GHz through 100 GHz
- **1550.52**—1550.52 nm, corresponds to 50 GHz

Copyright © 2013, Juniper Networks, Inc.

- **1550.92**—1550.92 nm, corresponds to 50 GHz through 100 GHz
- **1551.32**—1551.32 nm, corresponds to 50 GHz
- **1551.72**—1551.72 nm, corresponds to 50 GHz through 100 GHz
- **1552.12**—1552.12 nm, corresponds to 50 GHz
- **1552.52**—1552.52 nm, corresponds to 50 GHz through 100 GHz
- **1552.93**—1552.93 nm, corresponds to 50 GHz
- **1553.33**—1554.33 nm, corresponds to 50 GHz through 100 GHz
- **1553.73**—1554.73 nm, corresponds to 50 GHz
- **1554.13**—1554.13 nm, corresponds to 50 GHz through 100 GHz
- **1554.54**—1554.54 nm, corresponds to 50 GHz
- **1554.94**—1554.94 nm, corresponds to 50 GHz through 100 GHz
- **1555.34**—1555.34 nm, corresponds to 50 GHz
- **1555.75**—1555.75 nm, corresponds to 50 GHz through 100 GHz
- **1556.15**—1556.15 nm, corresponds to 50 GHz
- **1556.55**—1556.55 nm, corresponds to 50 GHz through 100 GHz
- **1556.96**—1556.96 nm, corresponds to 50 GHz
- **1557.36**—1557.36 nm, corresponds to 50 GHz through 100 GHz
- **1557.77**—1557.77 nm, corresponds to 50 GHz
- **1558.17**—1558.17 nm, corresponds to 50 GHz through 100 GHz
- **1558.58**—1558.58 nm, corresponds to 50 GHz
- **1558.98**—1558.98 nm, corresponds to 50 GHz through 100 GHz
- **1559.39**—1559.39 nm, corresponds to 50 GHz
- **1559.79**—1559.79 nm, corresponds to 50 GHz through 100 GHz
- **1560.20**—1560.20 nm, corresponds to 50 GHz
- **1560.61**—1560.61 nm, corresponds to 50 GHz through 100 GHz
- **1561.01**—1561.01 nm, corresponds to 50 GHz
- **1561.42**—1561.42 nm, corresponds to 50 GHz through 100 GHz
- **1561.83**—1561.83 nm, corresponds to 50 GHz
- **1562.23**—1562.23 nm, corresponds to 50 GHz through 100 GHz
- **1562.64**—1562.64 nm, corresponds to 50 GHz
- **1563.05**—1563.05 nm, corresponds to 50 GHz through 100 GHz
- **1563.45**—1563.45 nm, corresponds to 50 GHz

Copyright © 2013, Juniper Networks, Inc.

- **1563.86**—1563.86 nm, corresponds to 50 GHz through 100 GHz

Default: 1550.12—1550.12 nm, corresponds to 50 GHz through 100 GHz

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Ethernet DWDM Interface Wavelength Overview on page 441

west-interface

Syntax

```
west-interface {
  node-id mac-address;
  control-channel channel-name {
  interface-none
  ring-protection-link-end;
}
```

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Define one of the two interface ports for Ethernet ring protection, the other being defined by the **east-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

NOTE: **Always configure this port second, after configuring the east-interface statement.**

The statements are explained separately.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Ethernet Ring Protection Switching Overview on page 527
- *Ethernet Ring Protection Using Ring Instances for Load Balancing*
- *east-interface*
- **ethernet-ring on page 567**
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Configuring Ethernet Ring Protection Switching (CLI Procedure)*

working-circuit

Syntax working-circuit *group-name*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure the working router in an APS circuit pair.

Options ***group-name***—Circuit's group name.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Basic Automatic Protect Switching*
- *protect-circuit*

 Copyright © 2013, Juniper Networks, Inc.

PART 4

Troubleshooting

 Copyright © 2013, Juniper Networks, Inc.

CHAPTER 35

Investigate Fast Ethernet and Gigabit Ethernet Interfaces

Investigating Interface Steps and Commands

This section includes the following information to assist you when troubleshooting ATM interfaces:

Investigating Interface Steps and Commands Overview

The "Monitoring Interfaces" on page 699 section helps you determine the nature of the interface problem. The "Performing a Loopback Test on an Interface" on page 700 section provides information to help you isolate the source of the problem. The "Locating Interface Alarms" on page 702 section explains some of the alarms and errors for the media.

Monitoring Interfaces

Problem The following steps are a general outline of how you monitor interfaces to determine the nature of interface problems. For more detailed information on a specific interface, see the corresponding monitor interfaces section.

Solution To monitor interfaces, follow these steps:

1. Display the status of an interface.
2. Display the status of a specific interface.

Copyright © 2013, Juniper Networks, Inc.

3. Display extensive status information for a specific interface.
4. Monitor statistics for an interface.

The Table 38 on page 700 lists and describes the operational mode commands you use to monitor interfaces.

Table 38: Commands Used to Monitor Interfaces

CLI Command	Description
show interfaces terse ***interface-name*** For example: **show interfaces terse t1***	Displays summary information about the named interfaces.
show interfaces ***interface-name*** For example: **show interfaces t1-***x/y/z*	Displays static status information about a specific interface.
show interfaces ***interface-name*** **extensive** For example: **show interfaces t1-***x/y/z* **extensive**	Displays very detailed interface information about a specific interface.
monitor interface ***interface-name*** For example: **monitor interface t1-***x/y/z*	Displays real-time statistics about a physical interface, updated every second.

Performing a Loopback Test on an Interface

Problem The following steps are a general outline of how you use loopback testing to isolate the source of the interface problem. For more detailed information on a specific interface, see the corresponding loopback section.

Solution To use loopback testing for interfaces, follow these steps:

1. To diagnose a suspected hardware problem:
   a. Create a loopback.
   b. Set clocking to internal. (Not for Fast Ethernet/Gigabit Ethernet or Multichannel DS3 interfaces.)
   c. Verify that the status of the interface is up.
   d. Configure a static address resolution protocol table entry. (Fast Ethernet/Gigabit Ethernet interfaces only)
   e. Clear the interface statistics.
   f. Force the link layer to stay up.
   g. Verify the status of the logical interface.

 Copyright © 2013, Juniper Networks, Inc.

h. Ping the interface.

i. Check for interface error statistics.

2. To diagnose a suspected connection problem:

   a. Create a loop from the router to the network.

   b. Create a loop to the router from various points in the network.

The Table 39 on page 701 lists and describes the operational and configuration mode commands you use to perform loopback testing on interfaces (the commands are shown in the order in which you perform them).

Table 39: Commands Used to Perform Loopback Testing on Interfaces

CLI Statement or Command	Interface Type	Description
[edit interfaces *interface-name* interface-options] **set loopback (local | remote)**	All interfaces	The **loopback** statement at the hierarchy level configures a loopback on the interface. Packets can be looped on either the local router or the remote channel service unit (CSU). To turn off loopback, remove the **loopback** statement from the configuration.
show	All interfaces	Verify the configuration before you commit it.
commit	All interfaces	Save the set of changes to the database and cause the changes to take operational effect. Use after you have verified a configuration in all configuration steps.
[edit interfaces *interface-name*] **set clocking internal**	T1, T3, ATM, and SONET interfaces	The **clocking** statement at this hierarchy level configures the clock source of the interface to internal.
show interfaces *interface-name*	Used for all interfaces	Display static status information about a specific interface.
[edit interfaces *interface-name* unit *logical-unit-number* family inet address *ip-address*] **set arp *ip-address mac mac-address***	Fast Ethernet and Gigabit Ethernet interfaces	The **arp** statement at this hierarchy level defines mappings between IP and Media Access Control (MAC) addresses.
show arp no-resolve	Fast Ethernet and Gigabit Ethernet interfaces	Display the entries in the ARP table without attempting to determine the hostname that corresponds to the IP address (the **no-resolve** option).
clear interfaces statistics *interface-name*	All interfaces	Reset the statistics for an interface to zero.

Copyright © 2013, Juniper Networks, Inc.

Table 39: Commands Used to Perform Loopback Testing on Interfaces *(continued)*

CLI Statement or Command	Interface Type	Description
[edit interfaces *interface-name*] set encapsulation cisco-hdlc	T1, T3, SONET, and Multichannel DS3 interfaces	The **encapsulation** statement at this hierarchy level sets the encapsulation to the Cisco High-level Data-Link Control (HDLC) transport protocol on the physical interface.
[edit interfaces *interface-name*] set no-keepalives	T1, T3, SONET, and Multichannel DS3 interfaces	The **no-keepalives** statement at this level disables the sending of keepalives on the physical interface.
show interfaces *interface-name* terse	T1, T3, and SONET interfaces	Display summary information about interfaces. (Use to display the status of the logical interfaces for these interfaces.)
ping interface t1-*x/y/z local-ip-address* bypass-routing count 1000 rapid	All interfaces	Check the reachability of network hosts by sending **ICMP ECHO_REQUEST** messages to elicit **ICMP ECHO_RESPONSE** messages from the specified host. Use the **bypass-routing** option to ping a local system through an interface that has no route through it. The **count** option sends 1000 ping requests through the system. Type **Ctrl+C** to interrupt a **ping** command.
show interfaces *interface-name* extensive	All interfaces	Display very detailed interface information about a specific interface.

Locating Interface Alarms

Problem Locating alarms and errors for the media can be a simple process.

Solution To locate interface alarms and errors, use the **show interfaces *interface-name* extensive** command and examine the output for active alarms and defects.

Monitor Fast Ethernet and Gigabit Ethernet Interfaces

- Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces on page 702
- Monitor Fast Ethernet and Gigabit Ethernet Interfaces on page 703
- Fiber-Optic Ethernet Interface Specifications on page 710

Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Purpose To monitor Fast Ethernet and Gigabit Ethernet interfaces and begin the process of isolating interface problems when they occur.

 Copyright © 2013, Juniper Networks, Inc.

Action Table 40 on page 703 provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

Table 40: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
"Monitor Fast Ethernet and Gigabit Ethernet Interfaces" on page 703	
1. Display the Status of Fast Ethernet Interfaces on page 703	**show interfaces terse (fe* | ge*)**
2. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 705	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
3. Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface on page 706	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
4. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 709	**monitor interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
5. Fiber-Optic Ethernet Interface Specifications on page 710	

Meaning You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

Monitor Fast Ethernet and Gigabit Ethernet Interfaces

By monitoring Fast Ethernet and Gigabit Ethernet interfaces, you begin to isolate Fast Ethernet and Gigabit Ethernet interface problems when they occur.

To monitor your Fast Ethernet and Gigabit Ethernet interfaces, follow these steps:

1. Display the Status of Fast Ethernet Interfaces on page 703
2. Display the Status of Gigabit Ethernet Interfaces on page 704
3. Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 705
4. Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface on page 706
5. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 709

Display the Status of Fast Ethernet Interfaces

Purpose To display the status of Fast Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Action user@host> **show interfaces terse (fe* | ge*)**

Sample Output

```
user@host> show interfaces terse fe*
Interface       Admin Link Proto Local                 Remote
fe-2/1/0        up    up
```

```
fe-2/1/0.0      up    up   inet  10.116.115.217/29
fe-3/0/2        up    down
fe-3/0/2.0      up    down
fe-3/0/3        up    up
fe-3/0/3.0      up    up   inet  192.168.223.65/30
fe-4/1/0        down  up
fe-4/1/0.0      up    down inet  10.150.59.133/30
fe-4/1/1        up    up
fe-4/1/1.0      up    up   inet  10.150.59.129/30
fe-4/1/2        up    down
fe-4/1/2.0      up    down
```

Meaning The sample output lists only the Fast Ethernet interfaces. It shows the status of both the physical and logical interfaces. For a description of what the output means, see Table 41 on page 704.

Table 41: Status of Fast Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
fe-2/1/0 Admin Up Link Up	fe-2/1/0.0 Admin Up Link Up	This interface has both the physical and logical links up and running.
fe-3/0/2 Admin Up Link Down	fe-3/0/2.0 Admin Up Link Down	This interface has the physical link down, the link layer down, or both down (**Link Down**). The logical link is also down as a result.
fe-4/1/0 Admin Down Link Up	fe-4/1/0.0 Admin Up Link Down	This interface is administratively disabled and the physical link is healthy (**Link Up**), but the logical interface is not established. The logical interface is down because the physical link is disabled.
fe-4/1/2 Admin Up Link Down	fe-4/1/2.0 Admin Up Link Down	This interface has both the physical and logical links down.

Display the Status of Gigabit Ethernet Interfaces

Purpose To display the status of Gigabit Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Sample Output Action

```
user@host> show interfaces terse ge*
Interface       Admin Link Proto Local                 Remote
ge-2/2/0        down  down
ge-2/2/0.0      up    down inet  65.113.23.105/30
```

 Copyright © 2013, Juniper Networks, Inc.

```
ge-2/3/0          up    up
ge-2/3/0.0        up    up    inet   65.115.56.57/30
ge-3/1/0          up    up
ge-3/1/0.0        up    up    inet   65.115.56.193/30
ge-3/2/0          up    down
```

Meaning This sample output lists only the Gigabit Ethernet interfaces. It shows the status of both the physical and logical interfaces. See Table 42 on page 705 for a description of what the output means.

Table 42: Status of Gigabit Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
ge-2/2/0 Admin Down Link Down	ge-2/2/0.0 Admin Up Link Down	This interface is administratively disabled (**Admin Down**). Both the physical and logical links are down (**Link Down**).
ge-2/3/0 Admin Up Link Up	ge-2/3/0.0 Admin Up Link Up	This interface has both the physical and logical links up and running.
ge-3/2/0 Admin Up Link Down	ge-3/2/0.0 Admin Up Link Down	This interface has both the physical link and the logical interface down.

Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface

Purpose To display the status of a specific Fast Ethernet or Gigabit Ethernet interface when you need to investigate its status further, use the following Junos OS CLI operational mode command:

Action user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**

Sample Output 1

The following sample output is for a Fast Ethernet interface with the physical link up:

```
user@host> show interfaces fe-2/1/0
Physical interface: fe-2/1/0, Enabled, Physical link is Up
  Interface index: 31, SNMP ifIndex: 35
  Description: customer connection
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:86:71:1b, Hardware address: 00:90:69:86:71:1b
  Input rate     : 25768 bps (11 pps), Output rate: 1576 bps (3 pps)
  Active alarms  : None
```

Copyright © 2013, Juniper Networks, Inc.

```
  Active defects : None
  Logical interface fe-2/1/0.0 (Index 2) (SNMP ifIndex 43)
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.116.151.218/29, Local: 10.119.115.217
        Broadcast: 10.116.151.225
```

Sample Output 2

The following output is for a Gigabit Ethernet interface with the physical link up:

```
user@host> show interfaces ge-3/1/0
Physical interface: ge-3/1/0, Enabled, Physical link is Up
  Interface index: 41, SNMP ifIndex: 55
  Description: customer connection
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 1000mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:85:71:99, Hardware address: 00:90:69:85:71:99
  Input rate     : 7412216 bps (1614 pps), Output rate: 2431184 bps (1776 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-3/1/0.0 (Index 11) (SNMP ifIndex 57)
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.117.65.192/30, Local: 10.115.65.193
        Broadcast: 10.115.65.195
```

Meaning The first line of sample output 1 and 2 shows that the physical link is up. This means that the physical link is healthy and can pass packets. Further down the sample output, look for active alarms and defects. If you see active alarms or defects, to further diagnose the problem, see Step 3, "Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface" on page 706, to display more extensive information about the Fast Ethernet interface and the physical interface that is down.

Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface

Purpose To display extensive status information about a specific Fast Ethernet or Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**

Sample Output

The following sample output is for a Fast Ethernet interface:

```
user@router> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Up
  Interface index: 47, SNMP ifIndex: 38
  Description: Test
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
```

 Copyright © 2013, Juniper Networks, Inc.

```
Interface flags: SNMP-Traps
Link flags     : None
Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
Traffic statistics:
 Input  bytes  :           373012658                    0 bps
 Output bytes  :           153026154                 1392 bps
 Input  packets:             1362858                    0 pps
 Output packets:             1642918                    3 pps
 Input errors:
  Errors: 0 , Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
  L3 incompletes: 1 , L2 channel errors: 0 , L2 mismatch timeouts: 0
  FIFO errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
  HS link CRC errors: 0, FIFO errors: 0
Active alarms  : None
Active defects : None
MAC statistics:                      Receive        Transmit
  Total octets                     439703575       177452093
  Total packets                      1866532         1642916
  Unicast packets                     972137         1602563
  Broadcast packets                       30            2980
  Multicast packets                   894365           37373
  CRC/Align errors                         0               0
  FIFO errors                              0               0
  MAC control frames                       0               0
  MAC pause frames                         0               0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                 1866532
  Input packet rejects                     0
   Input DA rejects           503674
   Input SA rejects                0
  Output packet count                                1642916
  Output packet pad count                                  0
  Output packet error count                                0
  CAM destination filters: 5, CAM source filters: 0
 Autonegotiation information:
   Negotiation status: Complete, Link partner status: OK
  Link partner: Full-duplex, Flow control: None
PFE configuration:
  Destination slot: 1, Stream number: 15
  CoS transmit queue bandwidth:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  CoS weighted round-robin:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
  Description: Test
  Flags: SNMP-Traps, Encapsulation: ENET2
  Protocol inet, MTU: 1500, Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.115.107.192/29, Local: 10.115.107.193
      Broadcast: 10.115.107.199
```

Copyright © 2013, Juniper Networks, Inc.

Meaning The sample output shows where the errors might be occurring and includes autonegotiation information. See Table 43 on page 708 for a description of errors to look for.

Table 43: Errors to Look For

Error	Meaning
Policed discards	Discarded frames that were not recognized or were not of interest.
L2 channel errors	Packets for which the router could not find a valid logical interface. For example, the packet is for a virtual LAN (VLAN) that is not configured on the interface.
MTU	The maximum transmission unit (MTU) must match the interface of either the router at the remote end of the Fast Ethernet or Gigabit Ethernet link, or that of the switch.
Input DA rejects	Number of packets with a destination Media Access Control (MAC) address that is not on the accept list. It is normal to see this number increment.
Input SA rejects	Number of packets with a source MAC address that is not on the accept list. This number only increments when source MAC address filtering is configured.

If the physical link is down, look at the active alarms and defects for the Fast Ethernet or Gigabit Ethernet interface and diagnose the Fast Ethernet or Gigabit Ethernet media accordingly. See "Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters" on page 722 for an explanation of Fast Ethernet and Gigabit Ethernet alarms.

Table 44 on page 708 lists and describes some MAC statistics errors to look for.

Table 44: MAC Statistics Errors

Error	Meaning
CRC/Align errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad FCS with an integral number of octets (**FCS Error**) or a bad FCS with a non-integral number of octets (**Alignment Error**).
MAC control frames	The number of MAC control frames.
MAC pause frames	The number of MAC control frames with **pause** operational code.
Jabber frames	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Note that this definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragment frames	The total number of packets received that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error an alignment error. Note that it is entirely normal for fragment frames to increment because both runts (which are normal occurrences due to collisions) and noise hits are counted.

 Copyright © 2013, Juniper Networks, Inc.

Autonegotiation is the process that connected Ethernet interfaces use to communicate the information necessary to interoperate. Table 45 on page 709 explains the autonegotiation information of the **show interface *interface-name* extensive** command output.

Table 45: Autonegotiation Information

Autonegotiation Field Information	Explanation
Negotiation status: Incomplete	The **Negotiation status** field shows **Incomplete** when the Ethernet interface has the speed or link mode configured.
Negotiation status: No autonegotiation	The **Negotiation status** field shows **No autonegotiation** when the remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
Negotiation status: Complete **Link partner status: OK**	The **Negotiation status** field shows **Complete** and the **Link partner** field shows **OK** when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process completes successfully.
Link partner: Half-duplex	The **Link partner** field can be **Full-duplex** or **Half-duplex** depending on the capability of the attached Ethernet device.
Flow control: Symmetric/asymmetric	The **Flow control** field displays the types of flow control supported by the remote Ethernet device.

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface

Purpose To monitor statistics for a Fast Ethernet or Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action user@host> **monitor interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**

CAUTION: **We recommend that you use the monitor interface fe-*fpc/pic/port* or monitor interface ge-*fpc/pic/port* command only for diagnostic purposes. Do not leave these commands on during normal router operations because real-time monitoring of traffic consumes additional CPU and memory resources.**

Sample Output

The following sample output is for a Fast Ethernet interface:

```
user@host> monitor interface fe-2/1/0
Interface: fe-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:                                          Current Delta
  Input bytes:             282556864218 (14208 bps)              [40815]
  Output bytes:             42320313078 (384 bps)                  [890]
  Input packets:              739373897 (11 pps)                   [145]
  Output packets:             124798688 (1 pps)                     [14]
Error statistics:
```

```
    Input errors:                        0                               [0]
    Input drops:                         0                               [0]
    Input framing errors:                0                               [0]
    Policed discards:              6625892                               [6]
    L3 incompletes:                     75                               [0]
    L2 channel errors:                   0                               [0]
    L2 mismatch timeouts:                0                               [0]
    Carrier transitions:                 1                               [0]
    Output errors:                       0                               [0]
    Output drops:                        0                               [0]
    Aged packets:                        0                               [0]
  Active alarms : None
  Active defects: None
  Input MAC/Filter statistics:
    Unicast packets              464751787                             [154]
    Packet error count                   0                               [0]
```

Meaning Use the information from this command to help narrow down possible causes of an interface problem.

NOTE: If you are accessing the router from the console connection, make sure you set the CLI terminal type using the **set cli terminal** command.

The statistics in the second column are the cumulative statistics since the last time they were cleared using the **clear interfaces statistics** ***interface-name*** command. The statistics in the third column are the cumulative statistics since the **monitor interface** ***interface-name*** command was executed.

If the input errors are increasing, verify the following:

1. Check the cabling to the router and have the carrier verify the integrity of the line. To verify the integrity of the cabling, make sure that you have the correct cables for the interface port. Make sure you have single-mode fiber cable for a single-mode interface and multimode fiber cable for a multimode interface.
2. For a fiber-optic connection, measure the received light level at the receiver end and make sure that it is within the receiver specification of the Ethernet interface. See "Fiber-Optic Ethernet Interface Specifications" on page 710 for the fiber-optic Ethernet interface specifications.
3. Measure the transmit light level on the Tx port to verify that it is within specification. See "Fiber-Optic Ethernet Interface Specifications" on page 710 for the optical specifications.

Fiber-Optic Ethernet Interface Specifications

Table 46 on page 711 shows the specifications for fiber-optic interfaces for Juniper Networks routers.

 Copyright © 2013, Juniper Networks, Inc.

Table 46: Fiber-Optic Ethernet Interface Specifications

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
Gigabit Ethernet					
Duplex SC connector					
LH optical interface	49.5-mile 70-km reach on 8.2-micrometer SMF	1480 to 1580 nm	-3 to +2 dBm	-3 dBm	-23 dBm (BER 1012) for SMF
LX optical interface	6.2-mile 10-km reach on 9/125-micrometer SMF 1804.5-ft 550-m reach on 62.5/125- and 50/125-micrometer MMF	1270 to 1355 nm	-11 to -3 dBm	-3 dBm	-19 dBm
SX optical interface	656-ft 200-m reach on 62.5/125-micrometer MMF 1640-ft 500-m reach on 50/125-micrometer MMF	830 to 860 nm	-9.5 to -4 dBm	-3 dBm	-17 dBm
Fast Ethernet 8-Port					
FX optical interface with MT-RJ connector	1.24-mile 2-km reach on 62.5/125-micrometer MMF	1270 to 1380 nm	-20 to -14 dBm	-14 dBm	-34 dBm

Use Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Purpose To use loopback testing to isolate Fast Ethernet and Gigabit Ethernet interface problems.

Action Table 47 on page 712 provides links and commands for using loopback testing for Fast Ethernet and Gigabit Ethernet interfaces.

Table 47: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
“Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 712	
1. Create a Loopback on page 713	
a. Create a Physical Loopback for a Fiber-Optic Interface on page 713	Connect the transmit port to the receive port.
b. Create a Loopback Plug for an RJ-45 Ethernet Interface on page 713	Cross pin 1 (TX+) and pin 3 (RX+) together, and pin 2 (TX-) and pin 6 (RX-) together.
c. Configure a Local Loopback on page 714	**[edit interfaces *interface-name* (fastether-options | gigether-options)]** **set loopback** **show** **commit**
2. Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 715	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**
3. Configure a Static Address Resolution Protocol Table Entry on page 717	**show interfaces ge-*fpc/pic/port*** **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]** **set arp *ip-address* mac *mac-address* show** **commit** **run show arp no-resolve**
4. Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 718	**clear interfaces statistics fe-*fpc/pic/port* | ge-*fpc/pic/port***
5. Ping the Fast Ethernet or Gigabit Ethernet Interface on page 719	**ping *remote-IP-address* bypass-routing interface (fe-*fpc/pic/port* | ge-*fpc/pic/port* count 100 rapid**
6. Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 720	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
“Diagnose a Suspected Circuit Problem” on page 721	Perform Steps 2 through 8 from “Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 712.

Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface

Problem When you suspect a hardware problem, take the following steps to help verify if there is a problem.

Solution To diagnose a suspected hardware problem with the Ethernet interface, follow these steps:

 Copyright © 2013, Juniper Networks, Inc.

Create a Loopback

You can create a physical loopback or configure a local loopback to help diagnose a suspected hardware problem. Creating a physical loopback is recommended because it allows you to test and verify the transmit and receive ports. If a field engineer is not available to create the physical loopback, you can configure a local loopback for the interface. The local loopback creates a loopback internally in the Physical Interface Card (PIC).

Create a Physical Loopback for a Fiber-Optic Interface

Action

To create a physical loopback at the port, connect the transmit port to the receive port using a known good fiber cable.

NOTE: Make sure you use single-mode fiber for a single-mode port and multimode fiber for a multimode port.

Meaning

When you create and then test a physical loopback, you are testing the transmit and receive ports of the PIC. This action is recommended if a field engineer is available to create the physical loop as it provides a more complete test of the PIC.

Create a Loopback Plug for an RJ-45 Ethernet Interface

Action

To create a loopback plug, cross pin 1 (TX+) and pin 3 (RX+) together, and cross pin 2 (TX-) and pin 6 (RX-) together. You need the following equipment to create the loopback:

- A 6-inch long CAT5 cable
- An RJ-45 connector
- A crimping tool

Copyright © 2013, Juniper Networks, Inc.

Figure 40 on page 714 illustrates how to create a loopback plug for an RJ-45 Ethernet interface.

Figure 40: RJ-45 Ethernet Loopback Plug

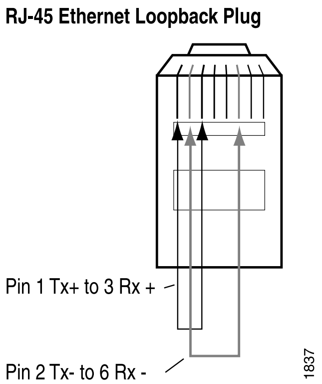

Meaning

When you create and then test a physical loopback, you are testing the RJ-45 interface of the PIC. This action is recommended if a field engineer is available to create the physical loop as it provides a more complete test of the PIC.

Configure a Local Loopback

Action

To configure a local loopback without physically connecting the transmit port to the receive port, follow these steps:

1. In configuration mode, go to the following hierarchy level:

   [edit]
   user@host# **edit interfaces** ***interface-name*** (fastether-options | gigether-options)

2. Configure the local loopback:

   [edit interfaces *interface-name* (fastether-options | gigether-options)]
   user@host# **set loopback**

3. Verify the configuration:

   user@host# **show**

   For example:

   [edit interfaces fe-1/0/0 fastether-options]
   user@host# **show**
   loopback;

4. Commit the change:

   user@host# **commit**

   For example:

   [edit interfaces fe-1/0/0 fastether-options]
   user@host# **commit**

 Copyright © 2013, Juniper Networks, Inc.

commit complete

When you create a local loopback, you create an internal loop on the interface being tested. A local loopback loops the traffic internally on that PIC. A local loopback tests the interconnection of the PIC but does not test the transmit and receive ports. On an Ethernet interface, you cannot create a remote loopback, therefore there is no option to use a **local** or **remote** statement. Simply including the **loopback** statement at the **[edit interfaces *interface-name* (fastether-options | gigether-options]** hierarchy level, places the interface into local loopback mode.

NOTE: Remember to delete the loopback statement after completing the test.

Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up

Purpose Display the status of the Fast Ethernet or Gigabit Ethernet interface to provide the information you need to determine whether the physical link is up or down.

Action To verify that the status of the Fast Ethernet or Gigabit Ethernet interface is up, use the following Junos OS command-line interface (CLI) operational mode command:

user@host> **show interfaces (fe-*fpc/port* | ge-*fpc/pic/port*)**

Sample Output

```
user@host# show interfaces fe-1/3/0
Physical interface: fe-1/3/0, Enabled, Physical link is Up
  Interface index: 44, SNMP ifIndex: 35
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps,  Loopback: Disabled , Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:db, Hardware address: 00:90:69:8d:2c:db
  Input rate     : 0 bps (0 pps), Output rate: 0 bps (0 pps)
  Active alarms : None
  Active defects : None
  MAC statistics:
    Input octets: 0, Input packets: 0, Output octets: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Autonegotiation information:
    Negotiation status: Incomplete, Link partner status: OK
    Link partner: Full-duplex, Flow control: None
```

Meaning

The sample output shows that the link is up and there are no alarms in this loopback configuration. When an internal loopback is configured, the physical loopback should come up without an alarm.

Copyright © 2013, Juniper Networks, Inc.

Sample Output

When you see that the physical link is down, there may be a problem with the port. The following output is an example of the show interfaces fe-*fpc/pic/port* command when the physical link is down:

```
user@router> show interfaces fe-1/3/0
Physical interface: fe-1/3/0, Enabled, Physical link is Down
  Interface index: 44, SNMP ifIndex: 35
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:db, Hardware address: 00:90:69:8d:2c:db
  Input rate     : 0 bps (0 pps), Output rate: 0 bps (0 pps)
  Active alarms : LINK
  Active defects : LINK
  MAC statistics:
    Input octets: 0, Input packets: 0, Output octets: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Autonegotiation information:
    Negotiation status: Incomplete, Link partner status: Down
    Reason: Link partner autonegotiation failure
    Link partner: Half-duplex, Flow control: None
```

Meaning The sample output shows that the physical link is down and there are active alarms and defects.

Table 48 on page 716 presents problem situations and actions for a physical link that is down.

Table 48: Problems and Solutions for a Physical Link That Is Down

Problem	Action
Cable mismatch	Verify that the fiber connection is correct.
Damaged and/or dirty cable	Verify that the fiber can successfully loop a known good port of the same type.
Too much or too little optical attenuation	Verify that the attenuation is correct per the PIC optical specifications.
The transmit port is not transmitting within the dBm optical range per the specifications	Verify that the Tx power of the optics is within range of the PIC optical specification.
Mismatch between the cable type and the port	Verify that a single-mode fiber cable is connected to a single-mode interface and that a multimode fiber cable is connected to a multimode interface. (This problem does not always cause the physical link to go down; errors and dropped packets are sometimes the result.)

 Copyright © 2013, Juniper Networks, Inc.

Configure a Static Address Resolution Protocol Table Entry

Purpose

Configure a static Address Resolution Protocol (ARP) entry to allow a packet to be sent out of a looped Ethernet interface.

NOTE: Remove the static ARP entry at the end of the loop test after you have completed the ping test, checked interface statistics, and monitored interface traffic.

Action

To configure a static ARP table entry for a Gigabit Ethernet interface, follow these steps. You can follow the same procedure to configure a static ARP entry for a Fast Ethernet interface.

1. Find the Media Access Control (MAC) address for the Gigabit Ethernet interface:

   user@host> **show interfaces ge-*fpc/pic/port***

2. In configuration mode, go to the following hierarchy level:

   [edit]
   user@host# **edit interfaces *interface-name* unit *logical-unit-number* family inet address *address***

3. Configure the static ARP entry:

   user@host# **set arp *ip-address* mac *mac-address***

   NOTE: The MAC address used should be the same as the physical address of the port being tested because this allows the port to receive the frames when you run the ping test.

4. Verify the configuration:

   user@host# **show**

5. Commit the configuration:

   user@host# **commit**

6. Verify that the static ARP entry is installed:

   user@host# **run show arp no-resolve**

Sample Output

```
user@host> show interfaces ge-7/2/1
Physical interface: ge-7/2/1, Enabled, Physical link is Down
  Interface index: 44, SNMP ifIndex: 35
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
```

```
  Interface flags: Hardware-Down SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:db, Hardware address: 00:90:69:8d:2c:db
  Input rate     : 0 bps (0 pps), Output rate: 0 bps (0 pps)
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host#  set arp 10.108.120.2 mac 00:90:69:8d:2c:db
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host# show
arp 10.108.120.2 mac 00:90:69:8d:2c:db;
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host# commit
commit complete
[edit interfaces ge-7/2/1 unit 0 family inet address 10.108.120.1/30]

user@host# run show arp no-resolve
MAC Address       Address         Interface     Flags
00:90:69:8d:2c:db 10.108.120.2    ge-7/2/1.0    permanent
00:e0:34:bb:8c:40 209.211.135.1   fxp0.0        none
00:a0:a5:28:0c:70 209.211.135.8   fxp0.0        none
00:a0:a5:12:12:c7 209.211.135.10  fxp0.0        none
00:90:ab:3c:68:a0 209.211.135.31  fxp0.0        none
08:00:20:a1:53:15 209.211.135.65  fxp0.0        none
00:a0:cc:66:3e:85 209.211.135.98  fxp0.0        none
Total entries: 7
```

Meaning

The sample output is for Step 1 through Step 6 and shows that a static ARP entry was configured on Gigabit Ethernet interface **ge-7/2/1**. The MAC address used is the same as the physical address of the port being tested because this allows the port to receive the frames when you run the ping test. The port is working as expected if you see that the time to live (TTL) expired; if you do not receive a response to your ping test, it indicates a hardware problem.

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics

Purpose

You must reset the Fast Ethernet and Gigabit Ethernet interface statistics before initiating the ping test. Resetting the statistics provides a clean start so that previous input/output errors and packet statistics do not interfere with the current diagnostics.

Action

To clear all statistics for the interface, use the following Junos OS CLI operational mode command:

user@host> **clear interfaces statistics (fe-*fpc/pic/port* | ge-*fpc/pic/port*)**

Sample Output

```
user@host>  clear interfaces statistics ge-7/2/0
user@host>
```

 Copyright © 2013, Juniper Networks, Inc.

Meaning

This command clears the interface statistics counters for the Gigabit Ethernet interface only.

Ping the Fast Ethernet or Gigabit Ethernet Interface

Purpose Use the ping command to verify the loopback connection.

Action To send ping packets from the Ethernet interface, use the following Junos OS CLI operational mode command:

user@host> **ping *remote-IP-address* bypass-routing interface (fe-*fpc/pic/port* | ge-*fpc/pic/port*) count 100 rapid**

Sample Output

```
user@router> ping 10.108.120.2 bypass-routing interface ge-7/2/1 count 100 rapid
PING 10.108.120.2 (10.108.120.2): 56 data bytes
36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e871   0 0000  01  01 cc5c 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e874   0 0000  01  01 cc59 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e878   0 0000  01  01 cc55 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e87c   0 0000  01  01 cc51 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e880   0 0000  01  01 cc4d 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e884   0 0000  01  01 cc49 10.108.120.1  10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
```

Meaning The sample output shows that the time to live (TTL) expired, indicating that the link is receiving the frames from the ping test. The MAC address used is the same as the physical address of the port being tested because this allows the port to accept the frames from the ping test. As the packet is looped over the link, you expect to receive a TLL exceeded message for each ping sent. These messages are generated because the ping packets are repeatedly looped between the router and the physical loopback. When the packet is sent to the other end of the link, which does not exist, the loopback returns the packet back to the same interface, where it is again subjected to the Packet Forwarding Engine fabric for routing. After the route lookup, the TTL is decremented, and the packet is again sent out of the looped interface. This process repeats until the packed is either lost, or the TLL expires with subsequent TTL expired message displayed. Should any errors occur, the packet is discarded and a time-out error is displayed, rather than the expected TTL expired message. Note that the default TTL for ICMP echo packets in Junos OS is 64. This means a given test packet must be successfully sent and received 63 times before a TTL expired message can be generated. You can alter the TTL value to adjust the tolerance for loss, for example, a value of 255 is the most demanding test because now the packet must be sent and received error free 254 times.

Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics

Purpose Persistent interface error statistics indicate that you need to open a case with the Juniper Networks Technical Assistance Center (JTAC).

Action To check the local interface for error statistics, use the following Junos OS CLI operational mode command:

user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**

Sample Output

```
user@router> show interfaces ge-7/2/1 extensive
Physical interface: ge-7/2/1, Enabled, Physical link is Up
  Interface index: 25, SNMP ifIndex: 32, Generation: 41
  Description: Test
  Link-level type: Ethernet, MTU: 4470, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:4c:17:b1, Hardware address: 00:90:69:4c:17:b1
  Statistics last cleared: 2002-01-07 17:53:19 UTC (2w2d 03:20 ago)
  Traffic statistics:
   Input  bytes  :        3799515503823                    0 bps
   Output bytes  :           7325566425                    0 bps
   Input  packets:           4628009535                    0 pps
   Output packets:             30678225                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0
  Output errors:
    Carrier transitions: 14, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0,
    FIFO errors: 0, HS link CRC errors: 0
  Active alarms : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                 3883579444813       7880356346
    Total packets                   4628009534         30678237
    Unicast packets                 4627879788         29893563
    Broadcast packets                       30              464
    Multicast packets                   129716           784210
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    VLAN tagged frames                       0
    Code violations                          0
  Filter statistics:
    Input packet count              4628009244
```

 Copyright © 2013, Juniper Networks, Inc.

```
      Input packet rejects                    0
      Input DA rejects                        0
      Input SA rejects                        0
      Output packet count                               30678237
      Output packet pad count                             856248
      Output packet error count                                0
      CAM destination filters: 9, CAM source filters: 0
    Autonegotiation information:
      Negotiation status: Complete, Link partner status: Ok, Link partner:
  Full-duplex,
      Flow control: None
    PFE configuration:
      Destination slot: 7
      CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                                %          bps   %        bytes
      0 best-effort             0            0   0            0       low    none
      1 expedited-forwarding    0            0   0            0       low    none
      2 assured-forwarding      0            0   0            0       low    none
      3 network-control         0            0   0            0       low    none
    Logical interface ge-7/2/1.0 (Index 23) (SNMP ifIndex 48) (Generation 38)
      Description: To Cosine Left 23/1
      Flags: SNMP-Traps Encapsulation: ENET2
      Protocol inet, MTU: 4456, Flags: None, Generation: 85 Route table: 0
        Addresses, Flags: Is-Preferred Is-Primary
         Destination: 10.108.120.0/30, Local: 10.108.120.1, Broadcast: 10.108.120.3,

          Generation: 81
      Protocol iso, MTU: 4453, Flags: None, Generation: 86 Route table: 0
```

Meaning Check for any error statistics. There should not be any input or output errors. If there are any persistent input or output errors, open a case with the Juniper Networks Technical Assistance Center (JTAC) at **support@juniper.net**, or at 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).

Diagnose a Suspected Circuit Problem

Purpose When you suspect a circuit problem, it is important to work with the transport-layer engineer to resolve the problem. The transport-layer engineer may create a loop to the router from various points in the network. You can then perform tests to verify the connection from the router to that loopback in the network.

Action After the transport-layer engineer has created the loop to the router from the network, you must verify the connection from the router to the loopback in the network. Follow Step 2 through Step 8 in "Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 712. Keep in mind that any problems encountered in the test indicate a problem with the connection from the router to the loopback in the network.

By performing tests to loopbacks at various points in the network, you can isolate the source of the problem.

Copyright © 2013, Juniper Networks, Inc.

Locate the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters

Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

Purpose To locate LINK alarm and major counters associated with Fast Ethernet and Gigabit Ethernet interfaces.

Action Table 49 on page 722 provides links and commands for locating LINK alarm and major counters for Fast Ethernet and Gigabit Ethernet interfaces.

Table 49: Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

Tasks	Command or Action
"Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm" on page 722	**show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**
"Fast Ethernet and Gigabit Ethernet Counters" on page 724	

Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm

Problem To display the Fast Ethernet or Gigabit Ethernet LINK alarm, use the following Junos OS command-line interface (CLI) operational mode command:

Solution user@host> **show interfaces (fe-*fpc/pic/port* | ge-*fpc/pic/port*) extensive**

Sample Output

The following sample output is for a Fast Ethernet interface:

```
user@host> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Down
  Interface index: 47, SNMP ifIndex: 38
  Description: Test
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
  Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
  Traffic statistics:
   Input  bytes  :           373012658                    0 bps
   Output bytes  :           153026154                 1392 bps
   Input  packets:             1362858                    0 pps
   Output packets:             1642918                    3 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
```

 Copyright © 2013, Juniper Networks, Inc.

```
  L3 incompletes: 1 , L2 channel errors: 0, L2 mismatch timeouts: 0
  FIFO errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
  HS link CRC errors: 0, FIFO errors: 0
```

Active alarms : LINK

```
Active defects : LINK
MAC statistics:                    Receive         Transmit
  Total octets                   439703575        177452093
  Total packets                    1866532          1642916
  Unicast packets                   972137          1602563
  Broadcast packets                     30             2980
  Multicast packets                 894365            37373
  CRC/Align errors                       0                0
  FIFO errors                            0                0
  MAC control frames                     0                0
  MAC pause frames                       0                0
  Oversized frames                       0
  Jabber frames                          0
  Fragment frames                        0
  VLAN tagged frames                     0
  Code violations                        0
Filter statistics:
  Input packet count               1866532
  Input packet rejects                   0
  Input DA rejects                  503674
  Input SA rejects                       0
  Output packet count                               1642916
  Output packet pad count                                 0
  Output packet error count                               0
  CAM destination filters: 5, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete, Link partner status: OK
  Link partner: Full-duplex, Flow control: None
PFE configuration:
  Destination slot: 1, Stream number: 15
  CoS transmit queue bandwidth:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  CoS weighted round-robin:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
  Description: Test
  Flags: SNMP-Traps, Encapsulation: ENET2
  Protocol inet, MTU: 1500, Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.115.107.192/29, Local: 10.115.107.193
      Broadcast: 10.115.107.199
```

Meaning

The sample output shows where the alarm and other errors might be occurring and any counters that are incrementing. The only alarm associated with Fast Ethernet or Gigabit Ethernet interfaces is the LINK alarm. A LINK alarm indicates a physical problem. To isolate where the physical problem might be occurring, conduct loopback testing. See “Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces” on page 711 for information on conducting a loopback test.

Copyright © 2013, Juniper Networks, Inc.

NOTE: Since link status is polled once every second, some items that require fast link down detection, such as Multiprotocol Label Switching (MPLS) fast reroute, take longer to execute.

Fast Ethernet and Gigabit Ethernet Counters

Problem

Table 50 on page 724 shows the major counters that appear in the output for the **show interfaces fe-*fpc/pic/port* extensive** and the **show interfaces ge-*fpc/pic/port* extensive** commands. These counters generally increment when there is a problem with a Fast Ethernet or Gigabit Ethernet interface. In the **Counters** column, the counters are listed in the order in which they are displayed in the output.

Solution

Table 50: Major Fast Ethernet and Gigabit Ethernet Counters

Counter	Description	Reason for Increment
Input Errors:		
Errors	The sum of the incoming frame aborts and frame check sequence (FCS) errors.	
Policed discards	The frames discarded by the incoming packet match code.	The frames were discarded because they were not recognized or of interest. Usually, this field reports protocols that the Junos OS does not handle.
Drops	The number of packets dropped by the output queue of the I/O Manager application-specific integrated circuit (ASIC).	If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism.
L3 incompletes	The number of packets discarded due to the packets failing Layer 3 header checks.	This counter increments when the incoming packet fails Layer 3 (usually IPv4) checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment.
L2 channel errors	The errors that occur when the software could not find a valid logical interface (such as **fe-1/2/3.0**) for an incoming frame.	This error increments when, for example, a lookup for a virtual LAN (VLAN) fails.
L2 mismatch timeouts	The count of malformed or short packets.	The malformed or short packets cause the incoming packet handler to discard the frame and be unreadable.
FIFO errors	The number of first in, first out (FIFO) errors in the receive direction as reported by the ASIC on the Physical Interface Card (PIC).	The value in this field should always be 0. If this value is not zero, cabling could be badly organized or the PIC could be broken.

Copyright © 2013, Juniper Networks, Inc.

Table 50: Major Fast Ethernet and Gigabit Ethernet Counters *(continued)*

Counter	Description	Reason for Increment
Output Errors		
Errors	The sum of outgoing frame aborts and FCS errors.	
Collisions	The number of Ethernet collisions.	The Fast Ethernet PIC supports only full-duplex operation, so this number should always remain 0. If it is incrementing, there is a software bug.
Drops	The 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 that is dropped by the ASIC's RED mechanism.
Aged packets	The number of packets that remained in shared packet SDRAM for so long that the system automatically purged them.	The value in this field should never increment. If it increments, it is probably a software bug or broken hardware.
HS link FCS errors, FIFO errors	The number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.	The value in this field should always be 0. If it increments, either the FPC or the PIC is broken.
Miscellaneous Counters		
Input DA rejects	The number of packets that the filter rejected because the destination Media Access Control (MAC) address of the packet is not on the accept list.	It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad Address Resolution Protocol (ARP) entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting).
Output packet pad count	The number of packets that the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware.	Usually, padding is done only on small ARP packets, but some very small Internet Protocol (IP) packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist, or it is misconfigured.
Output packet error count	Number of packets with an indicated error that the filter was given to transmit.	These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.
CAM destination filters, CAM source filters	The number of entries in the content-addressable memory (CAM) dedicated to destination and source MAC address filters.	There can be up to 64 source entries. If source filtering is disabled, which is the default, the value for these fields should be 0.

Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

PART 5

Index

Copyright © 2013, Juniper Networks, Inc.

Index

Symbols

A

B

C

 Copyright © 2013, Juniper Networks, Inc.

D

E

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.

T

U

V

 Copyright © 2013, Juniper Networks, Inc.

Index of Statements and Commands

 Copyright © 2013, Juniper Networks, Inc.

Copyright © 2013, Juniper Networks, Inc.