JUNIPER NETWORKS

Junos® OS

CoS on Ethernet IQ2 and Enhanced IQ2 PICs Feature Guide for Routing Devices

Release
14.1

Published: 2014-05-09

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

Junos® OS CoS on Ethernet IQ2 and Enhanced IQ2 PICs Feature Guide for Routing Devices
14.1

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- T Series
- M Series
- MX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

   ```
   commit {
     file ex-script-snippet.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 xi defines notice icons used in this guide.

Table 1: Notice Icons

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

Table 2 on page xii 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)	Encloses optional keywords or variables.	**stub <default-metric** *metric***>;**
| (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)	Encloses a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identifies 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

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

Convention	Description	Examples
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, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/

- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

CHAPTER 1

CoS on Ethernet IQ2 and Enhanced IQ2 PICs

CoS on Enhanced IQ2 PICs Overview

Some PICs, such as the Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Ethernet Enhanced IQ2 (IQ2E) PICs, have eight egress queues enabled by default on platforms that support eight queues.

The IQ2E PICs preserve all of the features of the IQ2 PICs, such as the default support for eight egress queues on platforms that support eight queues.

The IQ2E PICs add features such as the ability to perform hierarchical scheduling. You can mix IQ2 and IQ2E PICs on the same router.

The IQ2E PICs offer:

- Three levels of hierarchical CoS
- More granularity than a high priority queue
- 16,000 queues
- 2,000 schedulers with 8 queues
- 4,000 schedulers with 4 queues

The IQ2E PICs also offer automatic scheduler allocation across ports, so there is no need to reset the PIC when this changes. Random early detection (RED) keeps statistics on a per-drop-profile basis, improving the ability to perform network capacity planning.

When you include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level, each logical interface (unit) gets a dedicated scheduler (one scheduler is reserved for overflow). You can include the **per-session-scheduler** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level to shape Layer 2 Tunneling Protocol (L2TP) sessions. The behavior of these two-port scheduler modes is the same as in IQ2 PICs. However, IQ2E PICs use hierarchical schedulers and not shared schedulers; IQ2E PICs do not support the **shared-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level.

For more information about configuring hierarchical schedulers, including examples, see *Configuring Hierarchical Schedulers for CoS*.

You can shape traffic at the physical interface (port), logical interface (unit), or interface set (set of units) levels. Shaping is not supported at the queue level. However, you can include the **transmit-rate** statement with the **rate-limit** option at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level to police the traffic passing through a queue (but only in the egress direction). See "Configuring Rate Limits on IQ2 and Enhanced IQ2 PICs" on page 15.

At the physical interface (port) level, you can configure only a shaping rate (PIR). At the logical interface (unit) and interface set levels, you can configure both a shaping rate and a guaranteed rate (CIR). Note that the guaranteed rates at any level must be consistent with the parent level's capacity. In other words, the sum of the guaranteed rates on the logical interface (units) should be less than the guaranteed rate on the interface set, and the sum of the guaranteed rates on the logical interface (units) and interface sets should be less than the guaranteed rate on the physical interface (port).

You can control the rate of traffic that passes through the interface by configuring a policer overhead. When you configure a policer overhead, the configured policer overhead value is added to the length of the final Ethernet frame. This calculated length of the frame is used to determine the policer or the rate limit action. It does this because the policer overhead needs to be applied to policers just like shaping overhead is accounted for by shapers. The policer overhead is to be configured on the interface so that it is accounted for in the total packet length when policing traffic. See "Configuring a Policer Overhead" on page 24

The weighed RED (WRED) decision on the IQ2E PICs is done at the queue level. Once the accept or drop decision is made and the packet is queued, it is never dropped. Four drop profiles are associated with each queue: low, low-medium, medium-high, and high. WRED statistics are available for each loss priority (this feature is not supported on the IQ2 PICs). Also in contrast to the IQ2 PICs, the IQ2E PICs support WRED scaling profiles, allowing a single drop profile to be reused with a wide range of values. This practice increases the effective number of WRED drop profiles.

The IQ2E PICs provide four levels of strict priorities: strict-high, high, medium-high (medium-low) and low. In contrast to the IQ2 PICs, which support only one strict-high queue, the IQ2E PICs do not restrict the number of queues with a given priority. There is priority propagation among three levels: the logical interface, the logical interface set, and the physical port. These features are the same as those supported by Enhanced Queuing Dense Port Concentrators (DPCs) for Juniper Network MX Series 3D Universal

Edge Routers. For more information about configuring these features, see *Enhanced Queuing DPC Hardware Properties*.

The IQ2E PIC's queues are serviced with modified deficit round-robin (MDRR), as with the Enhanced Queuing DPCs. Excess bandwidth (bandwidth available after all guaranteed rates have been satisfied) can be shared equally or in proportion to the guaranteed rates. For more information about excess bandwidth sharing, see *Configuring Excess Bandwidth Sharing*.

Related Documentation

- egress-policer-overhead on page 39
- ingress-policer-overhead on page 42

CoS Capabilities and Limitations on IQ2 and IQ2E PICs (M Series and T Series Platforms)

This topic describes CoS scaling and performance parameters that apply to IQ2 and IQ2E PICs on M series and T series routers.

Rewrite Operations

802.1p or 802.1ad rewrite operations are done on the PIC. A total of eight rewrite markers of each type are supported on the PIC. For other rewrite operations, the numbers are the same as for any other M series FPCs. See *CoS Features and Limitations on M Series and T Series Routers*for details.

Classification

Classification is done on the PIC. There are eight classifier tables of each type (ieee-802.1p, mpls-exp, inet-precedence, dscp, and dscp-ipv6) supported per PIC.

For each classifier type, one table is reserved for a default classifier. This table is used when no classifier is configured, or when the number of tables configured exceeds eight.

Setting the Number of Egress Queues on IQ2 and Enhanced IQ2 PICs

Gigabit Ethernet IQ2 4-port and 8-port Type 2 PICs are oversubscribed, which means the amount of traffic coming to the PIC can be more than the maximum bandwidth from the PIC to the Flexible PIC Concentrator (FPC).

By default, PICs on M320, MX Series, and T Series routers support a maximum of four egress queues per interface. Some PICs, such as the IQ2 and IQ2E PICs, have eight egress queues enabled by default on platforms that support eight queues. You configure the number of egress queues as four or eight by including the **max-queues-per-interface** statement at the **[edit chassis fpc *slot-number* pic *pic-slot-number*]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-slot-number]
max-queues-per-interface (4 | 8);
```

The numerical value can be **4** or **8**.

For more information about configuring egress queues, see *Enabling Eight Queues on Interfaces*.

Shaping Granularity Values for Enhanced Queuing Hardware

Due to the limits placed on shaping thresholds used in the hierarchy, there is a granularity associated with the Enhanced IQ2 (IQ2E) PIC and the Enhanced Queuing (EQ) DPC. For these hardware models, the shaper accuracies differ at various levels of the hierarchy, with shapers at the logical interface level (Level 3) being more accurate than shapers at the interface set level (Level 2) or the port level (Level 1). Table 3 on page 6 shows the accuracy of the logical interface shaper at various rates for Ethernet ports operating at 1 Gbps.

Table 3: Shaper Accuracy of 1-Gbps Ethernet at the Logical Interface Level

Range of Logical Interface Shaper	Step Granularity
Up to 4.096 Mbps	16 Kbps
4.096 to 8.192 Mbps	32 Kbps
8.192 to 16.384 Mbps	64 Kbps
16.384 to 32.768 Mbps	128 Kbps
32.768 to 65.535 Mbps	256 Kbps
65.535 to 131.072 Mbps	512 Kbps
131.072 to 262.144 Mbps	1024 Kbps
262.144 to 1 Gbps	4096 Kbps

Table 4 on page 6 shows the accuracy of the logical interface shaper at various rates for Ethernet ports operating at 10 Gbps.

Table 4: Shaper Accuracy of 10-Gbps Ethernet at the Logical Interface Level

Range of Logical Interface Shaper	Step Granularity
Up to 10.24 Mbps	40 Kbps
10.24 to 20.48 Mbps	80 Kbps
10.48 to 40.96 Mbps	160 Kbps
40.96 to 81.92 Mbps	320 Kbps
81.92 to 163.84 Mbps	640 Kbps
163.84 to 327.68 Mbps	1280 Kbps

Table 4: Shaper Accuracy of 10-Gbps Ethernet at the Logical Interface Level *(continued)*

Range of Logical Interface Shaper	Step Granularity
327.68 to 655.36 Mbps	2560 Kbps
655.36 to 2611.2 Mbps	10240 Kbps
2611.2 to 5222.4 Mbps	20480 Kbps
5222.4 to 10 Gbps	40960 Kbps

Table 5 on page 7 shows the accuracy of the interface set shaper at various rates for Ethernet ports operating at 1 Gbps.

Table 5: Shaper Accuracy of 1-Gbps Ethernet at the Interface Set Level

Range of Interface Set Shaper	Step Granularity
Up to 20.48 Mbps	80 Kbps
20.48 Mbps to 81.92 Mbps	320 Kbps
81.92 Mbps to 327.68 Mbps	1.28 Mbps
327.68 Mbps to 1 Gbps	20.48 Mbps

Table 6 on page 7 shows the accuracy of the interface set shaper at various rates for Ethernet ports operating at 10 Gbps.

Table 6: Shaper Accuracy of 10-Gbps Ethernet at the Interface Set Level

Range of Interface Set Shaper	Step Granularity
Up to 128 Mbps	500 Kbps
128 Mbps to 512 Mbps	2 Mbps
512 Mbps to 2.048 Gbps	8 Mbps
2.048 Gbps to 10 Gbps	128 Mbps

Table 7 on page 7 shows the accuracy of the physical port shaper at various rates for Ethernet ports operating at 1 Gbps.

Table 7: Shaper Accuracy of 1-Gbps Ethernet at the Physical Port Level

Range of Physical Port Shaper	Step Granularity
Up to 64 Mbps	250 Kbps

Table 7: Shaper Accuracy of 1-Gbps Ethernet at the Physical Port Level *(continued)*

Range of Physical Port Shaper	Step Granularity
64 Mbps to 256 Mbps	1 Mbps
256 Mbps to 1 Gbps	4 Mbps

Table 8 on page 8 shows the accuracy of the physical port shaper at various rates for Ethernet ports operating at 10 Gbps.

Table 8: Shaper Accuracy of 10-Gbps Ethernet at the Physical Port Level

Range of Physical Port Shaper	Step Granularity
Up to 640 Mbps	2.5 Mbps
640 Mbps to 2.56 Gbps	10 Mbps
2.56 Gbps to 10 Gbps	40 Mbps

Understanding Burst Size Configuration on IQ2 and IQ2E Interfaces

You can explicitly configure the burst size for shapers in a traffic control profile for IQ2 and IQ2E interfaces. This feature is supported on M7i, M10i, M40e, M120, M320 routers and all T Series routers.

The shaping burst size determines the maximum number of bytes that can be sent through a shaper during a burst. The guaranteed burst size determines when the scheduler moves from green to yellow.

The burst size limits the number of credits that can be accumulated for scheduling. Configuring a burst size is only useful in the case when traffic is sent after a long lull period so that credits can be accumulated until the burst size limit is reached. When traffic is continuous, credits are not accumulated, and the burst size limit is not reached.

If no burst size value is specified when the shaping rate or guaranteed rate is configured, then a default burst size (expressed as a time value) is applied. The default shaping burst size is 10 ms of the shaping rate (that is, 10*shaping rate/1000 bytes). The minimum value is 2048 bytes to accommodate the minimum of 1 MTU.

The burst size value is adjusted and rounded off to meet the restrictions enforced by the hardware. Thus, the actual burst size in the hardware might vary slightly from the configured value.

To enable this feature, include the **burst-size** statement at the following hierarchy levels:

```
[edit class-of-service traffic-control-profiles shaping-rate]
[edit class-of-service traffic-control-profiles guaranteed-rate]
```

NOTE: The guaranteed-rate burst size value cannot be greater than the shaping-rate burst size.

Related Documentation

- Configuring Burst Size for Shapers on IQ2 and IQ2E Interfaces on page 16
- guaranteed-rate on page 41
- shaping-rate on page 56

Differences Between Gigabit Ethernet IQ and Gigabit Ethernet IQ2 PICs

Because Gigabit Ethernet IQ PICs and Gigabit Ethernet IQ2 PICs use different architectures, they differ in the following ways:

- Gigabit Ethernet IQ2 PICs support a transmission rate within a queue, but do not support an exact rate within a queue. You can apply a weight to a queue, but you cannot put an upper limit on the queue transmission rate that is less than the logical interface can support. Consequently, including the **exact** option with the **transmit-rate (*rate* | percent *percent*)** statement at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level is not supported for Gigabit Ethernet IQ2 interfaces.
- Gigabit Ethernet IQ2 PICs support only one queue in the scheduler map with **medium-high**, **high**, or **strict-high** priority. If more than one queue is configured with **medium-high**, **high**, or **strict-high** priority, the commit operation fails.
- To ensure that protocol control traffic (such as OSPF, BGP, and RIP) are not dropped at the oversubscribed ingress direction, the software puts control protocol packets into a separate control scheduler. There is one control scheduler per port. These control schedulers are implemented as strict-high priority, so they transmit traffic until they are empty.
- On Gigabit Ethernet IQ2 PICs, you can configure a single traffic-control profile to contain both a PIR (the **shaping-rate** statement) and a CIR (the **guaranteed-rate** statement). On Gigabit Ethernet IQ PICs, these statements are mutually exclusive.
- Gigabit Ethernet IQ2 PICs support only two fill levels in the RED drop profile. The recommended definition of the RED drop profile is as follows:

```
class-of-service {
  drop-profiles {
    drop-iq2-example1 {
      fill-level 20 drop-probability 0;
      fill-level 100 drop-probability 80;
    }
  }
}
```

  This configuration defines a drop profile with a linear drop probability curve when the fill level is between 20 and 100 percent, and a maximum drop probability of 80 percent.

You can configure more than two fill levels in a drop profile, but the software only uses the points (**min_fill_level**, **0**) and (**max_fill_level**, **max_probability**) and ignores other fill

levels. The drop probability at the minimum fill level is set to 0 percent even if you configure a non-zero drop probability value at the minimum fill level. The following example shows a sample configuration and the software implementation:

Configuration

```
class-of-service {
  drop-profiles {
    drop-iq2-example2 {
      fill-level 30 drop-probability 10;
      fill-level 40 drop-probability 20;
      fill-level 100 drop-probability 80;
    }
  }
}
```

Implementation

```
class-of-service {
  drop-profiles {
    drop-iq2-example2-implementation {
      fill-level 30 drop-probability 0;
      fill-level 100 drop-probability 80;
    }
  }
}
```

If you configure more than two fill levels, a system log message warns you that the software supports only two fill levels and displays the drop profile that is implemented.

Though the **interpolate** statement is supported in the definition of a RED drop profile, we do not recommend using it. The following example shows a sample configuration and the software implementation:

Configuration

```
class-of-service {
  drop-profiles {
    drop-iq2-example3 {
      interpolate {
        fill-level [ 30 50 80 ];
        drop-probability [ 10 20 40 ];
      }
    }
  }
}
```

When you use the **interpolate** statement and the maximum fill level is not 100 percent, the software adds the point (100, 100). Therefore, the drop-iq2-example3 drop profile is implemented as:

Implementation

```
class-of-service {
  drop-profiles {
    drop-iq2-example3-implementation {
      fill-level 2 drop-probability 0;
      fill-level 100 drop-probability 100;
    }
  }
}
```

The implemented minimum fill level is not 30 percent as configured, but 2 percent because of the 64-point interpolation.

CoS for L2TP Tunnels on Ethernet Interface Overview

For effective packet tunneling, CoS is implemented over L2TP tunnels. For Ethernet interfaces, CoS is supported for L2TP session traffic to a LAC on platforms configured as an LNS that include egress IQ2 or IQ2E PICs.

This feature is supported on the following platforms:

- M7i and M10i routers
- M120 routers

To enable session-aware CoS on an L2TP interface, include the **per-session-scheduler** statement at the **[edit interfaces unit *logical-unit-number*]** hierarchy level.

After CoS is configured on an L2TP tunnel, Junos OS dynamically creates a traffic shaper for the traffic-shaping-profile and the L2TP tunnel based on the tunnel identification number. This ensures that the packets are monitored at the LAC and classified to allow the traffic flow to be adjusted on congested networks.

This feature has the following limitations:

- Only 991 shapers are supported on each IQ2 or IQ2E PIC.
- For a 4-port IQ2E PIC, you can configure up to 1976 shapers for an 8-queue session and 3952 shapers for a 4-queue session.
- For an 8-port IQ2E PIC, you can configure up to 1912 shapers for an 8-queue session and up to 3824 shapers for a 4-queue session.
- Sessions in excess of the maximum supported values specified for the PICs cannot be shaped (but they can be policed).
- There is no support for PPP multilinks.
- The overall traffic rate cannot exceed the L2TP traffic rate, or else random drops result.
- There is no support for logical interface scheduling and shaping at the ingress because all schedulers are now reserved for L2TP.
- There is no support for physical interface rate shaping at the ingress.
- You cannot delete or deactivate the primary Ethernet interface on which the tunnel is established.

You can provide policing support for sessions with more than the maximum supported value on each IQ2 or IQ2E PIC. Each session can have four or eight different classes of traffic (queues). Each class needs its own policer; for example, one for voice and one for data traffic.

Related Documentation

- Configuring CoS for L2TP Tunnels on Ethernet Interfaces on page 25
- Configuring LNS CoS for Link Redundancy on page 26
- Example: L2TP LNS CoS Support for Link Redundancy on page 34

PART 2

Configuration

CHAPTER 2

Configuration Tasks

Configuring Rate Limits on IQ2 and Enhanced IQ2 PICs

You can rate-limit strict-high and high queues on IQ2 and IQ2E PICs. Without this limiting, traffic in higher priority queues can block the transmission of lower priority packets. Unless limited, higher priority traffic is always sent before lower priority traffic, causing the lower priority queues to "starve," which in turn leads to timeouts and unnecessary resending of packets.

On the IQ2 and IQ2E PICs you can rate-limit queues before the packets are queued for output. All packets exceeding the configured rate limit are dropped, so care is required when establishing this limit. For more information about configuring CoS on IQ2E PICs, see "CoS on Enhanced IQ2 PICs Overview" on page 3.

NOTE: **IQ2E PICs exclude the transmit rate of strict-high and high priority queues, thereby allowing low and medium priority queues to be configured up to 100 percent.**

To rate-limit queues, include the **transmit-rate** statement with the **rate-limit** option at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level:

```
[edit class-of-service schedulers scheduler-name]
transmit-rate rate rate-limit;
```

This example limits the transmit rate of a strict-high expedited-forwarding queue to 1 megabit per second (Mbps). The scheduler and scheduler map are defined and then applied to the traffic at the **[edit interfaces]** and **[edit class-of-service]** hierarchy levels:

```
[edit class-of-service]
schedulers {
   scheduler-1 {
      transmit-rate 1m rate-limit; # This establishes the limit
      priority strict-high;
   }
}
scheduler-maps {
   scheduler-map-1 {
      forwarding-class expedited-forwarding scheduler scheduler-1;
   }
}

[edit interfaces]
s0-2/1/0 {
   per-unit-scheduler;
   encapsulation frame-relay;
   unit 0 {
      dlci 1;
   }
}

[edit class-of-service]
interfaces {
   so-2/1/0 {
      unit 0 {
         scheduler-map scheduler-map-1;
         shaping-rate 2m;
      }
   }
}
```

You can issue the following operational mode commands to verify your configuration (the first shows the rate limit in effect):

- **show class-of-service scheduler-map** *scheduler-map-name*
- **show class-of-service interface** *interface-name*

Configuring Burst Size for Shapers on IQ2 and IQ2E Interfaces

This topic shows how to set the **burst-size** while configuring the **shaping-rate** and **guaranteed-rate** under the **[edit class-of-service traffic-control-profiles** *profile-name***]** hierarchy level.

In following configuration for tcp1, the **shaping-rate** burst size is set to 5 KB, and the **guaranteed-rate** burst size is set to 3 KB under the **traffic-control-profiles** statement. To apply this configuration to a logical interface (ifl), the traffic-control-profile is attached to the ifl.

```
class-of-service {
  traffic-control-profiles {
    tcp1 {
      shaping-rate 100m burst-size 5k;
      guaranteed-rate 50m burst-size 3k;
    }
    tcp2 {
      shaping-rate 100m burst-size 5k;
    }
  }
  interfaces {
    interface-set ifset1 {
      output-traffic-control-profile tcp1;
    }
    ge-1/2/1 {
      unit 0 {
        output-traffic-control-profile tcp1;
      }
    }
    ge-1/2/2 {
      output-traffic-control-profile tcp2;
    }
  }
}
```

Related Documentation

- Understanding Burst Size Configuration on IQ2 and IQ2E Interfaces on page 8
- guaranteed-rate on page 41
- shaping-rate on page 56

Configuring Shaping on 10-Gigabit Ethernet IQ2 PICs

The 10-Gigabit Ethernet IQ2 PIC (which has **xe-** interfaces) is unlike other Gigabit Ethernet IQ2 PICs in that it does not have oversubscription. The bandwidth from the PIC to the FPC is sufficient to transmit the full line rate. However, the 10-Gigabit Ethernet IQ2 PIC has the same hardware architecture as other Gigabit Ethernet IQ2 PICs and supports all the same class-of-service (CoS) features. For more information, see the PIC guide for your routing platform.

To handle oversubscribed traffic, you can configure input shaping and scheduling based on Layer 2, MPLS, and Layer 3 packet fields. Gigabit Ethernet IQ2 PICs also support simple filters, accounting, and policing. This chapter discusses input and output shaping and scheduling. For information about accounting and policing, see the *Junos OS Network Interfaces Library for Routing Devices*.

NOTE: The CoS functionality supported on Gigabit Ethernet IQ2 PICs is not available across aggregated Ethernet links. However, if you configure a CoS scheduler map on the link bundle, the configuration is honored by the individual links within that bundle.

Therefore, CoS behaves as configured on a per-link level, but not across the aggregated links. For example, if you configure a shaping transmit rate of 100 Mbps on an aggregated Ethernet bundle with three ports (by applying a scheduler for which the configuration includes the **transmit-rate** statement with the **exact** option at the **[edit class-of-service schedulers** *scheduler-name***]** hierarchy level), each port is provisioned with a 33.33 Mbps shaping transmit rate.

You can configure shaping for aggregated Ethernet interfaces that use interfaces originating from Gigabit Ethernet IQ2 PICs. However, you cannot enable shaping on aggregated Ethernet interfaces when the aggregate bundle combines ports from IQ and IQ2 PICs.

By default, transmission scheduling is not enabled on logical interfaces. Logical interfaces without shaping configured share a default scheduler. This scheduler has a committed information rate (CIR) that equals 0. (The CIR is the guaranteed rate.) The default scheduler has a peak information rate (PIR) that equals the physical interface shaping rate. The default operation can be changed by configuring the software.

NOTE: For Gigabit Ethernet IQ2 interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** logical interface counters. However, correct values display for both of these **Transit statistics** when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.

To configure input and output shaping and scheduling, include the following statements at the **[edit class-of-service]** and **[edit interfaces]** hierarchy levels of the configuration:

```
[edit class-of-service]
traffic-control-profiles profile-name {
  delay-buffer-rate (percent percentage | rate);
  excess-rate percent percentage;
  guaranteed-rate (percent percentage | rate);
  scheduler-map map-name;
  shaping-rate (percent percentage | rate);
}
interfaces {
  interface-name {
    input-scheduler-map map-name;
    input-shaping-rate rate;
```

```
    scheduler-map map-name; # Output scheduler map
    shaping-rate rate; # Output shaping rate
    }
    unit logical-unit-number {
      input-scheduler-map map-name;
      input-shaping-rate (percent percentage | rate);
      scheduler-map map-name;
      shaping-rate (percent percentage | rate);
      input-traffic-control-profile profile-name shared-instance instance-name;
      output-traffic-control-profile profile-name shared-instance instance-name;
    }
  }
}

[edit interfaces interface-name]
per-unit-scheduler;
shared-scheduler;
```

NOTE: As indicated by the preceding configuration, the **scheduler-map** and **shaping-rate** statements can be included at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level. However, we do not recommend this configuration. Include the **output-traffic-control-profile** statement instead.

Related Documentation

- *Simple Filter Overview*
- *How Simple Filters Evaluate Packets*
- *Guidelines for Configuring Simple Filters*
- *Guidelines for Applying Simple Filters*

Configuring Traffic Control Profiles for Shared Scheduling and Shaping

Shared scheduling and shaping allows you to allocate separate pools of shared resources to subsets of logical interfaces belonging to the same physical port. You configure this by first creating a traffic-control profile, which specifies a shaping rate and references a scheduler map. You must then share this set of shaping and scheduling resources by applying an instance of the traffic-control profile to a subset of logical interfaces. You can apply a separate instance of the same (or a different) traffic-control profile to another subset of logical interfaces, thereby allocating separate pools of shared resources.

To configure a traffic-control profile, perform the following steps:

1. Include the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   shaping-rate (percent percentage | rate);
   ```

   You can configure the shaping rate as a percentage from 1 through 100 or as an absolute rate from 1000 through 160,000,000,000 bits per second (bps). The shaping

rate corresponds to a peak information rate (PIR). For more information, see *Oversubscribing Interface Bandwidth*.

2. Include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   scheduler-map map-name;
   ```

   For information about configuring schedulers and scheduler maps, see *Configuring Schedulers* and *Configuring Scheduler Maps*. Gigabit Ethernet IQ2 interfaces support up to eight forwarding classes and queues.

3. Include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   delay-buffer-rate (percent percentage | rate);
   ```

   You can configure the delay-buffer rate as a percentage from 1 through 100 or as an absolute rate from 1000 through 160,000,000,000 bits per second. The delay-buffer rate controls latency. For more information, see *Oversubscribing Interface Bandwidth* and *Providing a Guaranteed Minimum Rate*.

4. Include the **guaranteed-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

   ```
   [edit class-of-service traffic-control-profiles profile-name]
   guaranteed-rate (percent percentage | rate);
   ```

   You can configure the guaranteed rate as a percentage from 1 through 100 or as an absolute rate from 1000 through 160,000,000,000 bps. The guaranteed rate corresponds to a committed information rate (CIR). For more information, see *Providing a Guaranteed Minimum Rate*.

   You must now share an instance of the traffic-control profile.

To share an instance of the traffic-control profile, perform the following steps:

1. Include the **shared-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit interfaces interface-name]
   shared-scheduler;
   ```

   This statement enables logical interfaces belonging to the same physical port to share one set of shaping and scheduling resources.

NOTE: On each physical interface, the shared-scheduler and per-unit-scheduler statements are mutually exclusive. Even so, you can configure one logical interface for each shared instance. This effectively provides the functionality of per-unit scheduling.

2. To apply the traffic-control profile to an input interface, include the **input-traffic-control-profile** and **shared-instance** statements at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

   ```
   [edit class-of-service interfaces interface-name unit logical-unit-number]
   input-traffic-control-profile profile-name shared-instance instance-name;
   ```

   These statements are explained in Step 3.

3. To apply the traffic-control profile to an output interface, include the **output-traffic-control-profile** and **shared-instance** statements at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

   ```
   [edit class-of-service interfaces interface-name unit logical-unit-number]
   output-traffic-control-profile profile-name shared-instance instance-name;
   ```

   The profile name references the traffic-control profile you configured in Step 1 through Step 4 of the "Configuring Traffic Control Profiles for Shared Scheduling and Shaping" section. The shared-instance name does not reference a configuration. It can be any text string you wish to apply to multiple logical interfaces that you want to share the set of resources configured in the traffic-control profile. Each logical interface shares a set of scheduling and shaping resources with other logical interfaces that are on the same physical port and that have the same shared-instance name applied.

   This concept is demonstrated in "Example: Configuring Shared Resources on Ethernet IQ2 Interfaces" on page 30.

NOTE: You cannot include the **output-traffic-control-profile** statement in the configuration if any of the following statements are included in the logical interface configuration: **scheduler-map**, **shaping-rate**, **adaptive-shaper**, or **virtual-channel-group** (the last two are valid on J Series routers only).

Related Documentation

- *Scheduling and Shaping in Hierarchical CoS Queues for Traffic Routed to GRE Tunnels*

Configuring a Separate Input Scheduler for Each Interface

As an alternative to shared input traffic-control profiles, you can configure each interface to use its own input scheduler. For each physical interface, you can apply an input scheduler map to the physical interface or its logical interfaces, but not both.

For information about configuring schedulers and scheduler maps, see *Configuring Schedulers* and *Configuring Scheduler Maps*. Gigabit Ethernet IQ2 interfaces support up to eight forwarding classes and queues.

To configure a separate input scheduler on the physical interface, include the **input-scheduler-map** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level:

```
[edit class-of-service interfaces interface-name]
input-scheduler-map map-name;
```

To configure a separate input scheduler on a logical interface, perform the following steps:

1. Include the **input-scheduler-map** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

   [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]
   input-scheduler-map *map-name*;

2. For the corresponding physical interface, you must also include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

   [edit interfaces *interface-name*]
   per-unit-scheduler;

The **per-unit-scheduler** statement enables one set of output queues for each logical interface configured under the physical interface.

On Gigabit Ethernet IQ2 PIC interfaces, configuration of the **per-unit-scheduler** statement requires that you configure VLAN tagging also. When you include the **per-unit-scheduler** statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a dual-port Gigabit Ethernet IQ PIC, the maximum number is 384.

Configuring Per-Unit Scheduling for GRE Tunnels Using IQ2 and IQ2E PICs

M7i, M10i, M120,T Series, and TX Matrix routers with Intelligent Queuing 2 (IQ2) PICs and Intelligent Queuing 2 Enhanced (IQ2E) PICs support per unit scheduling for Generic Routing Encapsulation (GRE) tunnels, adding all the functionality of tunnel PICs to GRE tunnels. The class of service (CoS) for the GRE tunnel traffic will be applied as the traffic is looped through the IQ2 and IQ2E PIC.

Shaping is performed on full packets that pass through the GRE tunnel.

IQ2 and IQ2E PICs support all interfaces that are supported on tunnel PICs, as follows:

- **gr-*fpc/pic/port***
- **vt-*fpc/pic/port***
- **lt-*fpc/pic/port***
- **ip-*fpc/pic/port***
- **pe-*fpc/pic/port***
- **pd-*fpc/pic/port***
- **mt-*fpc/pic/port***

The port variable is always zero.

The provided tunnel functionality is the same as that of regular tunnel PICs.

You can specify that IQ2 and IQ2E PICs work exclusively in tunnel mode or as a regular PIC. The default setting uses IQ2 and IQ2E PICs as a regular PIC. To configure exclusive

tunnel mode, use the **tunnel-only** statement at the **[edit chassis fpc *slot-number* pic *slot-number* tunnel-services]** hierarchy level.

You can use the **show interfaces queue gr-*fpc/pic/port*** command to display statistics for the specified tunnel.

IQ2E PIC schedulers can be dynamically allocated across ports.

When IQ2 and IQ2E PICs work exclusively as a tunnel PIC, they support the same number of tunnel logical interfaces as regular tunnel PICs; for example each PIC can support 4K **gr-** logical interfaces.

NOTE: This feature supports only **traffic-control-profile** on **gr-** logical interfaces. It does not support Class of Service (CoS) on **gr-** logical interfaces.

Also, a scheduler is allocated for a **gr-** logical interface only when there is a traffic-control profile configured for it.

The **tunnel-only** statement is used to specify that the IQ2 or IQ2E PIC needs to work in tunnel mode, as follows:

```
[edit]
chassis {
   fpc slot-number {
      pic slot-number {
         tunnel-services {
            tunnel-only;
         }
      }
   }
}
```

The PIC will be automatically bounced when the tunnel services configuration is changed.

The **chassis traffic-manager** mode must have the ingress traffic manager enabled in order for the tunnel-services to work correctly.

On **gr-** interfaces, you can configure the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces gr-*fpc/pic/port* unit *logical-unit-number*]** hierarchy level:

```
[edit]
class-of-service {
   traffic-control-profiles {
      tcp {
         shaping-rate rate;
      }
   }
   interfaces {
      gr-fpc/pic/port {
         unit logical-unit-number {
            output-traffic-control-profile tcp
         }
      }
   }
```

```
}
```

The **gr-** logical interfaces without an explicit CoS configuration are not assigned a dedicated scheduler. These use a reserved scheduler meant for all unshaped tunnel traffic; that is, all traffic on **gr-** logical interfaces that do not have CoS configured and all traffic from other types of tunnels.

For more information on **chassis tunnel-services** configuration, see the *Junos OS Administration Library for Routing Devices*.

To view the configuration and statistics for GRE tunnel logical interfaces, use the **show interfaces queue gr-** command.

Related Documentation

- *Junos OS Administration Library for Routing Devices*
- CLI Explorer
- CoS on Enhanced IQ2 PICs Overview on page 3

Configuring Hierarchical Input Shapers

You can apply input shaping rates to both the physical interface and its logical interfaces. The rate specified at the physical level is distributed among the logical interfaces based on their input shaping-rate ratio.

To configure an input shaper on the physical interface, include the **input-shaping-rate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level:

```
[edit class-of-service interfaces interface-name]
input-shaping-rate rate;
```

To configure an input shaper on the logical interface, include the **input-shaping-rate** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
input-shaping-rate (percent percentage | rate);
```

For each logical interface, you can specify a percentage of the physical rate or an actual rate. The software converts actual rates into percentages of the physical rate.

Configuring a Policer Overhead

Configuring a policer overhead allows you to control the rate of traffic sent or received on an interface. When you configure a policer overhead, the configured policer overhead value (bytes) is added to the length of the final Ethernet frame. This calculated length of frame is used to determine the policer or the rate limit action. Therefore, the policer overhead enables you to control the rate of traffic sent or received on an interface. You can configure the policer overhead to rate-limit queues and Layer 2 and MAC policers. The policer overhead and the shaping overhead can be configured simultaneously on an interface.

This feature is supported on M Series and T Series routers with IQ2 PICs or IQ2E PICs, and on MX Series DPCs.

To configure a policer overhead for controlling the rate of traffic sent or received on an interface:

1. In the **[edit chassis]** hierarchy level in configuration mode, create the interface on which to add the policer overhead to input or output traffic.

   ```
   [edit chassis]
   user@host# edit fpc fpc pic pic
   ```

   For example:

   ```
   [edit chassis]
   user@host# edit fpc 0 pic 1
   ```

2. Configure the policer overhead to control the input or output traffic on the interface. You could use either statement or both the statements for this configuration.

   ```
   [edit chassis fpc fpc pic pic]
   user@host# set ingress-policer-overhead bytes;
   user@host# set egress-policer-overhead bytes;
   ```

   For example:

   ```
   [edit chassis fpc 0 pic 1]
   user@host# set ingress-policer-overhead 10;
   user@host# set egress-policer-overhead 20;
   ```

3. Verify the configuration:

```
[edit chassis]
user@host# show
fpc 0 {
    pic 1 {
        ingress-policer-overhead 10;
        egress-policer-overhead 20;
    }
}
```

NOTE: When the configuration for the policer overhead bytes on a PIC is changed, the PIC goes offline and then comes back online. In addition, the configuration in the CLI is on a per-PIC basis and, therefore, applies to all the ports on the PIC.

Related Documentation

- egress-policer-overhead on page 39
- ingress-policer-overhead on page 42

Configuring CoS for L2TP Tunnels on Ethernet Interfaces

The Layer 2 Tunneling Protocol (L2TP) is often used to carry traffic securely between an L2TP Network Server (LNS) to an L2TP Access Concentrator (LAC). CoS is supported

for L2TP session traffic to a LAC on platforms configured as an LNS that include egress IQ2 and IQ2E Ethernet PICs. The following routers support this feature:

- M7i and M10i routers
- M120 routers

To configure CoS for L2TP on Ethernet interfaces:

1. Configure L2TP services on the Ethernet interface.
2. On the Ethernet interface, enable session-aware CoS for L2TP sessions.

   [[edit interfaces *interface-name* unit *logical-unit-number*]
   user@host# set per-session-scheduler

3. Configure the traffic manager in the IQ2 or IQ2E PIC to enable per-session CoS support.

   [edit chassis fpc *slot-number* pic *pic-number*]
   user@host# set traffic-manager mode-session-shaping

4. (Optional) To fine tune the system, you may also set the traffic-manager mode to session-shaping and configure the value of ingress-shaping-overhead parameter from 50 through 130 depending on your network requirement.

   [edit chassis fpc *slot-number* pic pic-number]
   user@host# set traffic-manager ingress-shaping-overhead value mode-session-shaping

NOTE: If you deactivate or delete the primary Ethernet interface on which the L2TP tunnel is configured, the tunnel with sessions having CoS is torn down.

After CoS is enabled for L2TP tunnels on Ethernet interface, you can run the **show class-of-service l2tp-session** command to verify the mapping of CoS with the configured L2TP session.

Related Documentation

- *L2TP Minimum Configuration*
- *Configuring CoS for L2TP Tunnels on ATM Interfaces*
- CoS for L2TP Tunnels on Ethernet Interface Overview on page 11
- Configuring LNS CoS for Link Redundancy on page 26
- Example: L2TP LNS CoS Support for Link Redundancy on page 34
- **show class-of-service l2tp-session on page 68**

Configuring LNS CoS for Link Redundancy

You can configure multiple ports on the same IQ2 and IQ2E PICs to support link redundancy for CoS on L2TP tunnels configured on an Ethernet interface. Link redundancy is useful when the active port is unavailable due to events such as:

- Disconnection of the cable

- Rebooting of the remote end system
- Traffic re-routing through a different port due to network conditions

When link redundancy is enabled in such scenarios, the L2TP tunnels and its session are maintained by switching traffic to another port configured on the same IQ2 or IQ2E PIC.

To configure multiple ports (IQ and IQ2PE PIC) on an Ethernet interface for redundancy with CoS, configure per-session-scheduler for all Ethernet ports:

```
user@host#edit interfaces ge-2/0/0 unit 0 per-session-scheduler
user@host#edit interfaces ge-2/0/1 unit 0 per-session-scheduler
```

You can similarly configure all the ports on the IQ2 or IQ2E PIC to support link redundancy for CoS on L2TP tunnels.

NOTE:

- If one or more redundancy ports is removed from the configuration, the tunnels established through those redundancy ports also go down.
- You must configure per-session-scheduler for all the ports that are to be used for redundancy. If you do not do so, new tunnels or sessions with CoS do not get established.

Related Documentation

- *per-session-scheduler*

CHAPTER 3

Examples

Example: Configuring a CIR and a PIR on Ethernet IQ2 Interfaces

On Gigabit Ethernet IQ2 interfaces, you can configure a CIR (guaranteed rate) and a PIR (shaping rate) on a single logical interface. The configured rates are gathered into a traffic control profile. If you configure a traffic control profile with a CIR (guaranteed rate) only, the PIR (shaping rate) is set to the physical interface (port) rate.

In the following example, logical unit **0** has a CIR equal to 30 Mbps and a PIR equal to 200 Mbps. Logical unit **1** has a PIR equal to 300 Mbps. Logical unit **2** has a CIR equal to 100 Mbps and a PIR that is unspecified. For logical unit **2**, the software causes the PIR to be 1 Gbps (equal to the physical interface rate) because the PIR must be equal to or greater than the CIR.

Excess bandwidth is the leftover bandwidth on the port after meeting all the guaranteed rate requirements of the logical interfaces. For each port, excess bandwidth is shared as follows:

- Proportional to the guaranteed rate—This method is used if you configure one or more logical interfaces on a port to have a guaranteed rate.
- Proportional to the shaping rate—This method is used if you configure none of the logical interfaces on a port to have a guaranteed rate.

In this example, bandwidth is shared proportionally to the guaranteed rate because at least one logical interface has a guaranteed rate.

```
class-of-service {
  traffic-control-profiles {
    profile1 {
      shaping-rate 200m;
      guaranteed-rate 30m;
      delay-buffer-rate 150m;
      scheduler-map sched-map;
    }
    profile2 {
```

```
        shaping-rate 300m;
        delay-buffer-rate 500k;
        scheduler-map sched-map;
      }
      profile3 {
        guaranteed-rate 100m;
        scheduler-map sched-map;
      }
    }
    interfaces {
      ge-3/0/0 {
        unit 0 {
          output-traffic-control-profile profile1;
        }
        unit 1 {
          output-traffic-control-profile profile2;
        }
        unit 2 {
          output-traffic-control-profile profile3;
        }
      }
    }
}
```

Example: Configuring Shared Resources on Ethernet IQ2 Interfaces

For input traffic on physical interface **ge-1/2/3**, logical interface units **1**, **2**, and **3** are sharing one set of scheduler-shaper resources, defined by traffic-control profile **s1**. Logical interface units **4**, **5**, and **6** are sharing another set of scheduler-shaper resources, defined by traffic-control profile **s1**.

For output traffic on physical interface **ge-1/2/3**, logical interface units **1**, **2**, and **3** are sharing one set of scheduler-shaper resources, defined by traffic-control profile **s2**. Logical interface units **4**, **5**, and **6** are sharing another set scheduler-shaper resources, defined by traffic-control profile **s2**.

For each physical interface, the **shared-instance** statement creates one set of resources to be shared among units **1**, **2**, and **3** and another set of resources to be shared among units **4**, **5**, and **6**. Input and output shaping rates are configured at the physical interface level, which demonstrates the hierarchical shaping capability of the Gigabit Ethernet IQ2 PIC.

```
[edit]
class-of-service {
  traffic-control-profiles {
    s1 {
      scheduler-map map1;
      shaping-rate 100k;
    }
    s2 {
      scheduler-map map1;
      shaping-rate 200k;
    }
  }
```

```
forwarding-classes { # Map one forwarding class to one queue.
  queue 0 fc-be;
  queue 1 fc-be1;
  queue 2 fc-ef;
  queue 3 fc-ef1;
  queue 4 fc-af11;
  queue 5 fc-af12;
  queue 6 fc-nc1;
  queue 7 fc-nc2;
}
classifiers { # Map 802.1p bits to forwarding-class and loss-priority.
  ieee-802.1 ieee-8021p-table {
    forwarding-class fc-nc2 {
      loss-priority low code-points [111];
    }
    forwarding-class fc-nc1 {
      loss-priority low code-points [110];
    }
    forwarding-class fc-af12 {
      loss-priority low code-points [101];
    }
    forwarding-class fc-af11 {
      loss-priority low code-points [100];
    }
    forwarding-class fc-ef1 {
      loss-priority low code-points [011];
    }
    forwarding-class fc-ef {
      loss-priority low code-points [010];
    }
    forwarding-class fc-be1 {
      loss-priority low code-points [001];
    }
    forwarding-class fc-be {
      loss-priority low code-points [000];
    }
  }
}
interfaces {
  ge-1/2/3 {
    input-shaping-rate 500m;
    shaping-rate 500m; # Output shaping rate
    unit 0 { # Apply behavior aggregate classifier to an interface.
      classifiers {
        ieee-802.1 ieee-8021p-table;
      }
    }
    unit 1 {
      input-traffic-control-profile s1 shared-instance 1;
      output-traffic-control-profile s2 shared-instance 1;
    }
    unit 2 {
      input-traffic-control-profile s1 shared-instance 1;
      output-traffic-control-profile s2 shared-instance 1;
    }
    unit 3 {
```

```
            input-traffic-control-profile s1 shared-instance 1;
            output-traffic-control-profile s2 shared-instance 1;
          }
          unit 4 {
            input-traffic-control-profile s1 shared-instance 2;
            output-traffic-control-profile s2 shared-instance 2;
          }
          unit 5 {
            input-traffic-control-profile s1 shared-instance 2;
            output-traffic-control-profile s2 shared-instance 2;
          }
          unit 6 {
            input-traffic-control-profile s1 shared-instance 2;
            output-traffic-control-profile s2 shared-instance 2;
          }
        }
      }
    }
```

Configuring a Simple Filter

Configure a simple filter that overrides the classification derived from the lookup of the Layer 2 fields.

```
firewall {
  family inet {
    simple-filter sf-1 {
      term 1 {
        source-address 172.16.0.0/16;
        destination-address 20.16.0.0/16;
        source-port 1024-9071;
      }
      then { # Action with term-1
      forwarding-class fc-be1;
      loss-priority high;
    }
    term 2 {
      source-address 173.16.0.0/16;
      destination-address 21.16.0.0/16;
    }
    then { # Action with term-2
    forwarding-class fc-ef1;
    loss-priority low;
  }
}
interfaces { # Apply the simple filter.
ge-1/2/3 {
  unit 0 {
    family inet {
      simple-filter {
        input sf-1;
      }
    }
  }
}

class-of-service {
  scheduler-maps { # Configure a custom scheduler map.
```

```
    map1 {
      forwarding-class fc-be scheduler sch-Q0;
      forwarding-class fc-be1 scheduler sch-Q1;
      forwarding-class fc-ef scheduler sch-Q2;
      forwarding-class fc-ef1 scheduler sch-Q3;
      forwarding-class fc-af11 scheduler sch-Q4;
      forwarding-class fc-af12 scheduler sch-Q5;
      forwarding-class fc-nc1 scheduler sch-Q6;
      forwarding-class fc-nc2 scheduler sch-Q7;
    }
}
schedulers { # Define schedulers.
sch-Q0 {
    transmit-rate percent 25;
    buffer-size percent 25;
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile drop-default;
}
sch-Q1 {
    transmit-rate percent 5;
    buffer-size temporal 2000;
    priority high;
    drop-profile-map loss-priority any protocol any drop-profile drop-ef;
}
sch-Q2 {
    transmit-rate percent 35;
    buffer-size percent 35;
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile drop-default;
}
sch-Q3 {
    transmit-rate percent 5;
    buffer-size percent 5;
    drop-profile-map loss-priority any protocol any drop-profile drop-default;
}
sch-Q4 {
    transmit-rate percent 5;
    priority high;
    drop-profile-map loss-priority any protocol any drop-profile drop-ef;
}
sch-Q5 {
    transmit-rate percent 10;
    priority high;
    drop-profile-map loss-priority any protocol any drop-profile drop-ef;
}
sch-Q6 {
    transmit-rate remainder;
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile drop-default;
}
sch-Q7 {
    transmit-rate percent 5;
    priority high;
    drop-profile-map loss-priority any protocol any drop-profile drop-default;
}
```

Example: L2TP LNS CoS Support for Link Redundancy

This example shows how link redundancy is supported when CoS for L2TP is configured on Ethernet interfaces.

NOTE: In this example, support for link redundancy is demonstrated by manually disabling the interface. However, link redundancy is also supported when the interface goes down due to events such as disconnection of the cable or rebooting of the remote end system.

Requirements

Before you begin:

- Configure service and loopback interfaces.
- Configure CoS for L2TP.

This feature applies to M Series Multiservice Edge Router running Junos OS Release 12.1 or later and EX Series switches.

Overview

Junos OS now supports link redundancy for CoS configured on an L2TP LNS. In this example, we verify that an L2TP tunnel does not go down when the Ethernet interface, through which the tunnels and its sessions with CoS are established, goes down.

Figure 1 on page 35 shows a sample scenario in which L2TP access concentrator (LAC) devices operate on one side of an L2TP tunnel. LAC devices are configured with the address range of 192.168.100.0 with a subnet mask of 24. The LAC devices are connected to two backbone routers, P1 and P2. These two routers, P1 and P2, are connected over two Gigabit Ethernet ports on a single Ethernet IQ2 PIC to an L2TP network server (LNS). The LNS device is a router running Junos OS that supports redundancy for terminating L2TP sessions configured with CoS parameters. The CoS settings are applied on the interfaces using a RADIUS server when the L2TP session is set up. One of the Gigabit Ethernet interfaces on the IQ2 PIC present on the LNS device, ge-0/3/1, is connected to P1, while the other interface, ge-0/3/3, is linked to P2. Such a method of connection enables the subscriber sessions that reach the LAC devices to be forwarded to one of the two ports of the IQ2 PIC on the LNS device.

Figure 1: Topology to Verify Link Redundancy Support for L2TP LNS CoS

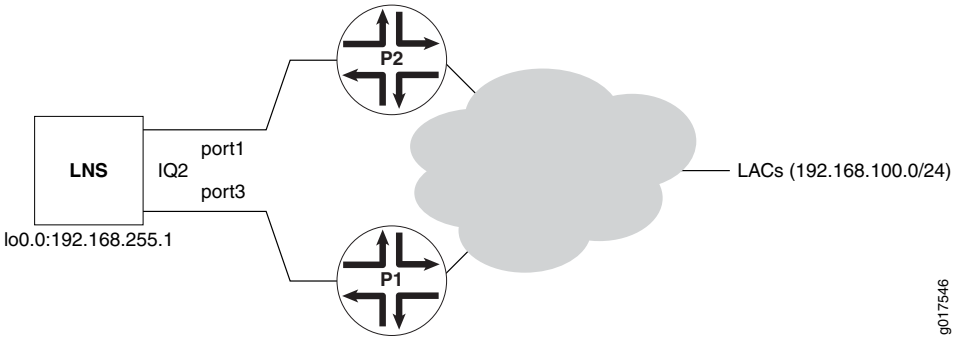

Configuration

Step-by-Step Procedure

To configure Ethernet interfaces for redundancy:

1. Configure Gigabit Ethernet interfaces.

```
[edit interfaces]
user@host# set ge-0/3/1 unit 0 family inet address 192.168.1.1/30
user@host# set ge-0/3/3 unit 0 family inet address 192.168.1.5/30
user@host# set ge-0/3/1 unit 0 per-session-scheduler
user@host# set ge-0/3/3 unit 0 per-session-scheduler
```

2. Configure static routing options.

```
[edit routing-options]
user@host# set static route 192.168.100.0/24 next-hop [ 192.168.1.2 192.168.1.6 ]
```

Step-by-Step Procedure

Verify that CoS is now implemented over L2TP on an Ethernet interface and the LAC is reachable.

1. Verify that LAC is reachable.

```
user@host> show route 192.168.100.1
inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.100.0/24   *[Static/5] 1d 02:09:09
                      to 192.168.1.2 via ge-0/3/1.0
                    > to 192.168.1.6 via ge-0/3/3.0
```

2. Bring up an L2TP session and verify that L2TP sessions come up.

```
user@host> show services l2tp session
Interface: sp-1/3/0, Tunnel group: GEN-TUN-GRP-BIO, Tunnel local ID: 44806
  Local Remote Interface State              Bundle Username
  ID    ID     unit
  12491  33795         1 Established             - test1
```

3. Send a traffic stream towards the subscriber.

4. Verify that the shaping at the subscriber end is as per the shaping rate configured.

```
user@host# show class-of-service l2tp-session
L2TP Session Username: test1, Index: 12491
Physical interface: ge-0/3/3, Index: 131
Queues supported: 4, Queues in use: 4
```

```
Scheduler map: GEN-SCHED-MAP-EF-65%, Index: 5212
Shaping rate: 2162200 bps
Encapsulation Overhead: 6, Cell Overhead: Enabled
```

In the output of the **show class-of-service l2tp-session** command, ge-0/3/3, index 131 represents the port used to establish the L2TP tunnel to which the current L2TP session belongs. It does not represent the port that was active when the L2TP session came up.

Verification

Verify that, when CoS is configured on an L2TP tunnel, link redundancy works if one of the ports on which the L2TP tunnel is established goes down.

- Bring Down ge-0/3/3 Interface Through Which the L2TP Tunnel Is Established on page 36
- Verify LAC Reachability and the Status of L2TP Sessions on page 36

Bring Down ge-0/3/3 Interface Through Which the L2TP Tunnel Is Established

Purpose Bring down the interface through which the L2TP session and its tunnels are established.

Action

```
[edit interfaces]
user@host# set ge-0/3/3 disable
user@host# commit
```

Verify LAC Reachability and the Status of L2TP Sessions

Purpose Verify that link redundancy works and the L2TP session does not go down when the active port on the IQ2 PIC is down. Verify that the traffic flow is unaffected after it is switched to another port configured on the same IQ2 or IQ2E PIC.

Action

```
user@host> show route 192.168.100.1
inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.100.0/24   *[Static/5] 1d 02:35:09
                      to 192.168.1.2 via ge-0/3/1.0


user@host> show services l2tp session
Interface: sp-1/3/0, Tunnel group: GEN-TUN-GRP-BIO, Tunnel local ID: 44806
  Local Remote Interface State               Bundle Username
  ID    ID     unit
  12491  33795         1 Established              - test1
```

Related Documentation

- Configuring LNS CoS for Link Redundancy on page 26

CHAPTER 4

Configuration Statements

delay-buffer-rate

Syntax delay-buffer-rate (percent *percentage* | *rate*);

Hierarchy Level [edit class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, base the delay-buffer calculation on a delay-buffer rate.

Default If you do not include this statement, the delay-buffer calculation is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured. For more information, see *Oversubscribing Interface Bandwidth*.

Options

percent *percentage*—For LSQ interfaces, delay-buffer rate as a percentage of the available interface bandwidth.

Range: 1 through 100 percent

rate—For IQ and IQ2 interfaces, delay-buffer rate, in bits per second (bps). You can specify a value 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).

Range: 1000 through 160,000,000,000 bps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Oversubscribing Interface Bandwidth*
- *Providing a Guaranteed Minimum Rate*
- Configuring Traffic Control Profiles for Shared Scheduling and Shaping on page 19
- output-traffic-control-profile on page 50

egress-policer-overhead

Syntax egress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1.

Description Add the specified number of bytes to the actual length of an Ethernet frame when determining the actions of Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card. You can configure egress policer overhead to account for egress *shaping* overhead bytes added to output traffic on the line card.

On M Series and T Series routers, this statement is supported on Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs. On MX Series routers, this statement is supported for interfaces configured on Dense Port Concentrators (DPCs).

NOTE: This statement is not supported on Modular Interface Cards (MICs) or Modular Port Concentrators (MPCs) in MX Series routers.

Options ***bytes***—Number of bytes added to a packet exiting an interface.

Range: 0–255 bytes

Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *egress-shaping-overhead*
- *Policer Overhead to Account for Rate Shaping Overview*
- *Example: Configuring Policer Overhead to Account for Rate Shaping*
- Configuring a Policer Overhead on page 24
- CoS on Enhanced IQ2 PICs Overview on page 3

excess-rate

Syntax excess-rate (percent *percentage* | proportion *value*);

Hierarchy Level [edit class-of-service schedulers *scheduler-name*],
[edit class-of-service traffic-control-profiles *traffic-control-profile-name*]

Release Information Statement introduced in Junos OS Release 9.3.
Application to the Multiservices PIC added in Junos OS Release 9.5.
Application to the MIC and MPC interfaces added in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.1X48R2 for PTX Series Packet Transport Routers.

Description For an Enhanced IQ PIC interfaces, Multiservices PIC interfaces, or MX Series router interfaces on MPCs or MICs, and T4000 router interfaces on Type 5 FPCs and EX Series switches, determine the percentage or proportion of excess bandwidth traffic to share.

NOTE: The proportion option provides a greater range of values over the percent option and hence influences the priorities assigned to the queues.

Options ***percentage***—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

Default: Excess bandwidth is shared in proportion to the configured transmit rate of each queue.

value—(M Series, MX Series, T Series routers and EX Series switches only) Proportion of the excess bandwidth to share. Option available at the **[edit class-of-service traffic-class-profiles *traffic-control-profile-name*]** hierarchy level only.

Range: 0 through 1000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Scheduler Transmission Rate*
- *Configuring Excess Bandwidth Sharing on IQE PICs*
- *Allocating Excess Bandwidth Among Frame Relay DLCIs on Multiservices PICs*
- *Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs*

guaranteed-rate

Syntax guaranteed-rate (percent *percentage* | *rate*) <burst-size *bytes*>;

Hierarchy Level [edit class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 7.6.
Option **burst-size** introduced for Enhanced Queuing (EQ) DPC interfaces in Junos OS Release 9.4.
Option **burst-size** introduced for MIC and MPC interfaces in Junos OS Release 11.4.
Option **burst-size** introduced for IQ2 and IQ2E interfaces in Junos OS Release 12.3

Description For Gigabit Ethernet IQ, Channelized IQ PICs, AS PIC FRF.16 LSQ interfaces, and EQ DPCs only, configure a guaranteed minimum rate. You can also configure an optional burst size for a logical interface on EQ DPCs and on IQ2 and IQ2E PICs. This can help to ensure that higher priority services do not starve lower priority services.

Default If you do not include this statement and you do not include the **delay-buffer-rate** statement, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 2 MTU-sized packets.

Options

percent *percentage*—For LSQ interfaces, guaranteed rate as a percentage of the available interface bandwidth.

Range: 1 through 100 percent

rate—For IQ and IQ2 interfaces, guaranteed rate, in bits per second (bps). You can specify a value 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).

Range: 1000 through 160,000,000,000 bps

burst-size *bytes*—(Optional) Maximum burst size, in bytes.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Providing a Guaranteed Minimum Rate*
- Configuring Traffic Control Profiles for Shared Scheduling and Shaping on page 19
- output-traffic-control-profile on page 50

ingress-policer-overhead

Syntax ingress-policer-overhead *bytes*;

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*]

Release Information Statement introduced before Junos OS Release 11.1

Description Add the configured number of bytes to the length of a packet entering the interface.

Options ***bytes***—Number of bytes added to a packet entering an interface.
Range: 0–255 bytes
Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *ingress-shaping-overhead*
- *Policer Overhead to Account for Rate Shaping Overview*
- *Example: Configuring Policer Overhead to Account for Rate Shaping*
- Configuring a Policer Overhead on page 24
- CoS on Enhanced IQ2 PICs Overview on page 3

input-scheduler-map

Syntax input-scheduler-map *map-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 7.6.

Description Associate a scheduler map with a physical or logical input interface. The **input-scheduler-map** and **input-traffic-control-profile** statements are mutually exclusive at the same hierarchy level.

input-scheduler-map is supported on the following Ethernet interfaces:

- IQ2 and IQ2E PICs
- DPCs and MPCs that support Enhanced Queuing (Q/EQ)
- MX80 with support for per-VLAN queuing

Options *map-name*—Name of scheduler map that you define at the **[edit class-of-service scheduler-maps]** hierarchy level.

default—The default scheduler mapping.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring a Separate Input Scheduler for Each Interface on page 21
- *Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs*
- input-traffic-control-profile on page 46

input-shaping-rate (Logical Interface)

Syntax input-shaping-rate (percent *percentage* | *rate*);

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ2, Enhanced Queuing DPC, MIC, and MPC interfaces, configure input traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. You can configure hierarchical shaping, meaning you can apply an input shaping rate to both the physical interface and the logical interface.

Default If you do not include this statement, logical interfaces share a default scheduler. This scheduler has a committed information rate (CIR) that equals 0. (The CIR is the guaranteed rate.) The default scheduler has a peak information rate (PIR) that equals the physical interface shaping rate.

Options **percent *percentage***—Shaping rate as a percentage of the available interface bandwidth.

Range: 0 through 100 percent

rate—Peak rate, in bits per second (bps). You can specify a value 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).

Range: 1000 through 160,000,000,000 bps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Ingress Hierarchical CoS on MIC and MPC Interfaces*
- Configuring Hierarchical Input Shapers on page 24
- *Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs*
- input-traffic-control-profile on page 46

input-shaping-rate (Physical Interface)

Syntax input-shaping-rate *rate*;

Hierarchy Level [edit class-of-service interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ2, Enhanced Queuing DPC, MIC, and MPC interfaces, configure input traffic shaping by specifying the amount of bandwidth to be allocated to the physical interface. You can configure hierarchical shaping, meaning you can apply an input shaping rate to both the physical interface and the logical interface.

Options *rate*—Peak rate, in bits per second (bps). You can specify a value 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).

Range: 1000 through 160,000,000,000 bps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Hierarchical Input Shapers on page 24
- *Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs*
- input-traffic-control-profile on page 46

input-traffic-control-profile

Syntax input-traffic-control-profile *profile-name* shared-instance *instance-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ2 and IQ2E PIC, Enhanced Queuing DPC , MIC, and MPC interfaces, apply an input traffic scheduling and shaping profile to the logical interface. The **input-traffic-control-profile** and **input-scheduler-map** statements are mutually exclusive at the same hierarchy level.

NOTE: The **shared-instance** statement applies only to Gigabit Ethernet IQ2 and IQ2E PICs.

Options ***profile-name***—Name of the traffic-control profile to be applied to this interface.

instance-name—Name of the shared scheduler and shaper to be applied to this interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Traffic Control Profiles for Shared Scheduling and Shaping on page 19
- *Configuring Ingress Hierarchical CoS on Enhanced Queuing DPCs*
- input-shaping-rate (Logical Interface) on page 44
- input-scheduler-map on page 43
- traffic-control-profiles on page 58

interfaces

Syntax

```
interfaces {
  interface-name {
    classifiers{
      dscp(classifier-name | default) {
    }
      ieee-802.1 (classifier-name | default)  vlan-tag (inner | outer | classifier-name);
      inet-precedence (rewrite-name | default);
    }
    input-scheduler-map map-name;
    input-shaping-rate rate;
    irb {
      unit logical-unit-number {
        classifiers {
          type (classifier-name | default);
        }
        rewrite-rules {
          dscp (rewrite-name | default);
          dscp-ipv6 (rewrite-name | default);
          exp (rewrite-name | default)protocol protocol-types;
          ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
          inet-precedence (rewrite-name | default);
        }
      }
    }
    member-link-scheduler (replicate | scale);
    rewrite-rules {
      dscp (rewrite-name | default);
      ieee-802.1 (rewrite-name | default) vlan-tag (outer);
      inet-precedence (rewrite-name | default);
      }
    }
    scheduler-map map-name;
    scheduler-map-chassis map-name;
    shaping-rate rate;
    unit logical-unit-number {
      classifiers {
        type (classifier-name | default) family (mpls | inet);
      }
      forwarding-class class-name;
      fragmentation-map map-name;
      input-shaping-rate (percent percentage | rate);
      input-traffic-control-profile profile-name shared-instance instance-name;
      output-traffic-control-profile profile-name shared-instance instance-name;
      per-session-scheduler;
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        exp (rewrite-name | default) protocol protocol-types;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default);
```

```
        }
        scheduler-map map-name;
        shaping-rate rate;
        translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
          | to-inet-precedence-from-inet-precedence) table-name;
      }
    }
    interface-set interface-set-name {
      excess-bandwidth-share;
      internal-node;
      output-traffic-control-profile profile-name;
      output-traffic-control-profile-remaining profile-name;
    }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.

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

- *Overview of BA Classifier Types*
- *Configuring Rewrite Rules*

max-queues-per-interface

Syntax max-queues-per-interface (8 | 4);

Hierarchy Level [edit chassis fpc *slot-number* pic *pic-number*],
[edit chassis lcc *number* fpc *slot-number* pic *pic-number*] (Routing Matrix)

Release Information Statement introduced before Junos OS Release 7.4.
Support for TX Matrix and TX Matrix Plus added in Junos OS Release 9.6.
On MIC or MPC interfaces on MX Series routers, configure eight egress queues.

Description On IQ interfaces on M320, T320, T640, TX Matrix, and TX Matrix Plus routers, or on MIC or MPC interfaces on MX Series routers, configure eight egress queues.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Junos OS to Support Eight Queues on IQ Interfaces for T Series and M320 Routers*
- *Configuring Up to 16 Forwarding Classes*
- *Enabling Eight Queues on ATM Interfaces*
- *Configuring the Maximum Number of Queues for Trio MPC/MIC Interfaces*

output-traffic-control-profile

Syntax output-traffic-control-profile *profile-name* shared-instance *instance-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*],
[edit class-of-service interfaces *interface-name* interface-set *interface-set-name*]

Release Information Statement introduced in Junos OS Release 7.6.
interface-set option added for Enhanced Queuing DPCs on MX Series routers in Junos OS Release 8.5.
interface-set option added for MIC and MPC interfaces on MX Series routers in Junos OS Release 10.2.
Support on GRE tunnel interfaces configured on physical and logical interfaces on MICs or MPCs in MX Series routers added in Junos OS Release 13.3.

Description Apply the specified CoS traffic control profile (traffic scheduling and shaping configuration objects) to the output traffic at the physical interface, logical interface, or interface set.

The statement is supported on the following interfaces:

- Channelized IQ PIC interfaces
- Gigabit Ethernet IQ, Gigabit Ethernet IQ2, and IQ2E PIC interfaces
- Link services IQ (LSQ) interfaces on AS PICs
- Enhanced Queuing DPC, MIC, and MPC interfaces on MX Series routers
- GRE tunnel interfaces configured on physical or logical interfaces hosted on MIC or MPC line cards in MX Series routers.

NOTE: **Interface sets (sets of interfaces used to configure hierarchical CoS schedulers on supported Ethernet interfaces) are not supported on GRE tunnel interfaces.**

The **shared-instance** statement is supported on Gigabit Ethernet IQ2 PICs only.

Options ***profile-name***—Name of the traffic-control profile to be applied to this interface

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Oversubscribing Interface Bandwidth*
- Configuring Traffic Control Profiles for Shared Scheduling and Shaping on page 19
- *Configuring Hierarchical Schedulers for CoS* (Enhanced Queuing DPC, MIC, and MPC interfaces on MX Series routers)

- *Configuring Interface Sets* (Enhanced Queuing DPC, MIC, and MPC interfaces on MX Series routers)
- *output-traffic-control-profile-remaining*
- traffic-control-profiles on page 58

scheduler-map (Interfaces and Traffic-Control Profiles)

Syntax scheduler-map *map-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*],
[edit class-of-service traffic-control-profiles]

Release Information Statement introduced before Junos OS Release 7.4.

Description For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, associate a scheduler map name with an interface or with a traffic-control profile.

For channelized OC12 intelligent queuing (IQ), channelized T3 IQ, channelized E1 IQ, and Gigabit Ethernet IQ interfaces only, you can associate a scheduler map name with a logical interface.

Options *map-name*—Name of the scheduler map.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Configuring Schedulers*
- *Oversubscribing Interface Bandwidth*
- output-traffic-control-profile on page 50

schedulers (Class of Service)

Syntax

```
schedulers {
  scheduler-name {
    adjust-minimum rate;
    adjust-percent percentage;
    buffer-size (seconds | percent percentage | remainder | temporal microseconds);
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
      (any | non-tcp | tcp) drop-profile profile-name;
    excess-priority [ low | medium-low | medium-high | high | none];
    excess-rate (percent percentage | proportion value);
    priority priority-level;
    shaping-rate (percent percentage | rate);
    transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series routers.

Description Specify the scheduler name and parameter values.

Options ***scheduler-name***—Name of the scheduler to be configured.

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

- *Schedulers Overview*
- *Default Schedulers Overview*
- *Configuring Schedulers*
- *Configuring a Scheduler*

shaping-rate (Applying to an Interface)

Syntax shaping-rate *rate*;

Hierarchy Level [edit class-of-service interfaces *interface-name*],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
[edit class-of-service interfaces *interface-name*] hierarchy level added in Junos OS Release 7.5.
Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers.

Description For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

For physical interfaces on IQ PICs and T4000 routers with Type 5 FPCs only, configure traffic shaping based on the rate-limited bandwidth of the total interface bandwidth.

Logical and physical interface traffic shaping rates are mutually exclusive. This means you can include the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level or the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, but not both.

NOTE: For MX Series routers and for EX Series switches, the shaping rate value for the physical interface at the **[edit class-of-service interfaces *interface-name*]** hierarchy level must be a minimum of 160 Kbps. If the value is less than the sum of the logical interface guaranteed rates, the user is not allowed to apply the shaping rate to a physical interface.

For T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of **shaping-rate** is limited by the maximum transmission rate of the interface.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate, as described in *Oversubscribing Interface Bandwidth*.

For FRF.15 and FRF.16 bundles on link services interfaces, only shaping rates based on percentage are supported.

Default If you do not include this statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level, the default physical interface bandwidth is the

average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options

rate—Peak rate, in bits per second (bps). You can specify a value 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).

Range: For logical interfaces, 1000 through 32,000,000,000 bps. For physical interfaces, 1000 through 160,000,000,000 bps.

NOTE: For all MX Series and EX series interfaces, the rate can be from 65,535 through 160,000,000,000 bps.

NOTE: For T4000 physical interfaces, the rate can be from 1000 through 160,000,000,000 bps.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Applying Scheduler Maps Overview*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Packet Transport Routers*

shaping-rate (Limiting Excess Bandwidth Usage)

Syntax shaping-rate (percent *percentage* | *rate*) <burst-size *bytes*>;

Hierarchy Level [edit class-of-service schedulers *scheduler-name*]

Release Information Statement introduced before Junos OS Release 7.4.
The **burst-size** option added for MIC and MPC interfaces on MX Series routers in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.2 for ACX Series Routers.

Description Define a limit on excess bandwidth usage for J Series routers and for MIC and MPC interfaces on MX Series routers.

The **transmit-rate** statement at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level configures the minimum bandwidth allocated to a queue. The transmission bandwidth can be configured as an exact value or allowed to exceed the configured rate if additional bandwidth is available from other queues. For J Series routers only, you limit the excess bandwidth usage with this statement.

You should configure the shaping rate as an absolute maximum usage and not the additional usage beyond the configured transmit rate.

Default If you do not include this statement, the default shaping rate is 100 percent, which is the same as no shaping at all.

Options **percent *percentage***—Shaping rate as a percentage of the available interface bandwidth.

Range: 0 through 100 percent

rate—Peak rate, in bits per second (bps). You can specify a value 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).

Range: 3200 through 32,000,000,000 bps

burst-size *bytes*—Maximum burst size, in bytes. The burst value determines the number of rate credits that can accrue when the queue or scheduler node is held in the inactive round robin.

Range: 0 through 1,000,000,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Applying Scheduler Maps Overview*

shaping-rate (Oversubscribing an Interface)

Syntax shaping-rate (percent *percentage* | *rate*) <burst-size *bytes*>;

Hierarchy Level [edit class-of-service traffic-control-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 7.6.
Option **burst-size** introduced for Enhanced Queuing (EQ) DPC interfaces on MX Series routers in Junos OS Release 9.4.
Option **burst-size** option introduced for MIC and MPC interfaces on MX Series routers in Junos OS Release 11.4.
Option **burst-size** introduced for IQ2 and IQ2E interfaces in Junos OS Release 12.3.

Description For Gigabit Ethernet IQ, Channelized IQ PIC, FRF.15 and FRF.16 LSQ interfaces, and for EQ DPC, MIC, and MPC interfaces on MX Series routers, configure a shaping rate for a logical interface. You can also configure an optional burst size for a logical interface on EQ DPC interfaces and on IQ2 and IQ2E PIC interfaces. This can help to ensure that higher-priority services do not starve lower-priority services.

For physical interfaces on T4000 router interfaces on Type 5 FPCs, configure traffic shaping rate.

The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).

Default The default behavior depends on various factors. For more information, see *Oversubscribing Interface Bandwidth*.

Options **percent *percentage***—For LSQ interfaces, shaping rate as a percentage of the available interface bandwidth.

Range: 1 through 100 percent

rate—For IQ and IQ2 interfaces, and T4000 routers with Type 5 FPCs, peak rate, in bits per second (bps). You can specify a value 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).

Range: IQ and IQ2 interfaces—1000 through 160,000,000,000 bps

Range: T4000 routers with Type 5 FPCs—the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of **shaping-rate** is limited by the maximum transmission rate of the interface.

burst-size *bytes*—(Optional) Maximum burst size, in bytes.

Range: 0 through 1,000,000,000

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Traffic Control Profiles for Shared Scheduling and Shaping on page 19
- *Oversubscribing Interface Bandwidth*
- output-traffic-control-profile on page 50

shared-instance

Syntax shared-instance *instance-name*;

Hierarchy Level [edit class-of-service interfaces *interface-name* unit *logical-unit-number* input-traffic-control-profile],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number* output-traffic-control-profile]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ2 and IQ2E PICs only, apply a shared traffic scheduling and shaping profile to the logical interface.

Options ***instance-name***—Name of the shared scheduler and shaper to be applied to this interface

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Shaping on 10-Gigabit Ethernet IQ2 PICs on page 17
- traffic-control-profiles on page 58

traffic-control-profiles

Syntax

```
traffic-control-profiles profile-name {
  adjust-minimum rate;
  atm-service (cbr | rtvbr | nrtvbr);
  delay-buffer-rate (percent percentage | rate);
  excess-rate (percent percentage  | proportion value );
  excess-rate-high (percent percentage | proportion value);
  excess-rate-low (percent percentage | proportion value);
  guaranteed-rate (percent percentage | rate) <burst-size bytes>;
  max-burst-size cells;
  overhead-accounting (frame-mode | cell-mode | frame-mode-bytes | cell-mode-bytes)
    <bytes (byte-value)>;
  peak-rate rate;
  scheduler-map map-name;
  shaping-rate (percent percentage | rate) <burst-size bytes>;
  shaping-rate-excess-high rate [ burst-size bytes ];
  shaping-rate-excess-low rate [ burst-size bytes ];
  shaping-rate-priority-high rate [ burst-size bytes ];
  shaping-rate-priority-low rate [ burst-size bytes ];
  shaping-rate-priority-medium rate [ burst-size bytes ];
  strict-priority-scheduler;
  sustained-rate rate;
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ, Channelized IQ PICs, FRF.15 and FRF.16 LSQ interfaces, Enhanced Queuing (EQ) DPCs, and PTX Series routers only, configure traffic shaping and scheduling profiles. For Enhanced EQ PICs, EQ DPCs, and PTX Series routers only, you can include the **excess-rate** statement.

Options ***profile-name***—Name of the traffic-control profile.

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

- *Oversubscribing Interface Bandwidth*
- *Understanding Scheduling on PTX Series Routers*
- output-traffic-control-profile on page 50

transmit-rate (Schedulers)

Syntax transmit-rate (*rate* | percent *percentage* | remainder) <exact | rate-limit>;

Hierarchy Level [edit class-of-service schedulers *scheduler-name*]

Release Information Statement introduced before Junos OS Release 7.4.
rate-limit option introduced in Junos OS Release 8.3. Applied to the Multiservices PICs in Junos OS Release 9.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for ACX Series Routers.

Description Specify the transmit rate or percentage for a scheduler.

Default If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.

Options

exact—(Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. This value should never exceed the rate-controlled amount. For PTX Series Packet Transport Routers, this option is allowed only on the non-strict-high (high, medium-high, medium-low, or low) queues.

percent *percentage*—Percentage of transmission capacity. A percentage of zero drops all packets in the queue.

Range: 0 through 100 percent for M, MX and T Series routers and EX Series switches; 1 through 100 percent for PTX Series Packet Transport Routers; 0 through 200 percent for the SONET/SDH OC48/STM16 IQE PIC

NOTE:

- On M Series Multiservice Edge Routers, for interfaces configured on 4-port E1 and 4-port T1 PICs only, you can configure a ***percentage*** value only from 11 through 100. These two PICs do not support transmission rates less than 11 percent.
- The configuration of the **transmit-rate percent 0 exact** statement at the **[edit class-of-service schedules *scheduler-name*]** hierarchy is ineffective on T4000 routers with Type 5 FPC.
- On MIC and MPC interfaces on MX Series routers, when the transmit rate is configured as a percentage and **exact** or **rate-limit** is enabled on a queue, the shaping rate of the parent node is used to compute the transmit rate. If **exact** or **rate-limit** is not configured, the guaranteed rate of the parent node is used to compute the transmit rate.

rate—Transmission rate, in bps. You can specify a value 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).

Range: 3200 through 160,000,000,000 bps

NOTE: For all MX Series interfaces, the rate can be from 65,535 through 160,000,000,000 bps.

rate-limit—(Optional) Limit the transmission rate to the rate-controlled amount by applying a policing action to the queue. Packets are hard-dropped when traffic exceeds the specified maximum transmission rate.

NOTE: For PTX Series Packet Transport Routers, this option is allowed only on the strict-high queue. We recommend that you configure rate limit on strict-high queues because the other queues may not meet their guaranteed bandwidths. The rate-limit option cannot rate limit the queue if strict-priority scheduling is configured with the *strict-priority-scheduler* statement.

NOTE: The configuration of the rate-limit statement is supported on T4000 routers only with a Type 5 FPC.

remainder—Use the remaining rate available.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Schedulers*
- *Configuring Scheduler Transmission Rate*
- *Understanding Scheduling on PTX Series Routers*

unit

Syntax

```
unit logical-unit-number {
  classifiers {
    type (classifier-name | default) family (mpls | all);
  }
  forwarding-class class-name;
  fragmentation-map map-name;
  input-traffic-control-profile profile-name shared-instance instance-name;
  output-traffic-control-profile profile-name shared-instance instance-name;
  per-session-scheduler;
  rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    exp (rewrite-name | default)protocol protocol-types;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
  }
  scheduler-map map-name;
  shaping-rate rate;
}
```

Hierarchy Level [edit class-of-service interfaces *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 16,384

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

- *Overview of BA Classifier Types*
- *Configuring Rewrite Rules*

per-unit-scheduler

Syntax per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 13.2 on 16x10GE MPC and MPC3E line cards.
Statement introduced in Junos OS Release 13.2 on PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 13.3 on MPC4E line cards.

Description For Channelized OC3 IQ, Channelized OC12 IQ, Channelized STM1 IQ, Channelized T3 IQ, Channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), link services (ls-) on J Series routers, Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-, 40-, and 100-Gigabit Ethernet interfaces (including the 16x10GE MPC), enable the association of scheduler map names with logical interfaces.

NOTE: Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.

NOTE: On Gigabit Ethernet IQ2 and IQ2-E PICs without the **per-unit-scheduler** statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the **per-unit-scheduler** statement, the entire PIC supports 1024 – 2 * number of ports (1024 minus two times the number of ports), because each port is allocated two default schedulers.

When including the **per-unit-scheduler** statement, you must also include the **vlan-tagging** statement or the **flexible-vlan-tagging** statement (to apply scheduling to VLANs) or the **encapsulation frame-relay** statement (to apply scheduling to DLCIs) at the **[edit interfaces *interface-name*]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*
- *vlan-tagging*
- *flexible-vlan-tagging*
- *Applying Scheduling and Shaping to VLANs*
- *Configuring Virtual LAN Queuing and Shaping on PTX Series Packet Transport Routers*

shared-scheduler

Syntax shared-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 7.6.

Description For Gigabit Ethernet IQ2 PICs only, enable shared schedulers and shapers on this interface. This statement and the **per-unit-scheduler** statement are mutually exclusive. Even so, you can configure one logical interface for each shared instance. This effectively provides the functionality of per-unit scheduling.

For Gigabit Ethernet IQ2 and Ethernet Enhanced IQ2 (IQ2E) PICs on M320 routers, enable shared schedulers on aggregated Ethernet interfaces in link protection mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Shaping on 10-Gigabit Ethernet IQ2 PICs on page 17
- traffic-control-profiles on page 58

PART 3

Administration

CHAPTER 5

Operational Mode Commands

- show class-of-service l2tp-session

show class-of-service l2tp-session

Syntax show class-of-service l2tp-session *session-id*

Release Information Command introduced in Junos OS Release 8.2.

Description Display CoS objects associated with an L2TP session on M7i, M10i, and M120 routers.

Options ***session-id***—L2TP session number for which you want to display a summary of CoS attributes.

Required Privilege Level view

List of Sample Output show class-of-service l2tp-session on page 68

Output Fields Table 9 on page 68 lists the output fields for the **show class-of-service l2tp-session** command. Output fields are listed in the approximate order in which they appear.

Table 9: show class-of-service l2tp-session Output Fields

Field Name	Field Description
L2TP Session Username	Username associated with the L2TP session.
Index	Session index identification number.
Physical interface	Physical interface on which the tunnel session is established.
Index	Index ID associated with the physical interface on which the tunnel session is established.
Queues supported	Number of scheduler queues supported for the L2TP session.
Queues in use	Number of scheduler queues active on the L2TP session.
Scheduler map	Scheduler map name associated with the session.
Index	Scheduler map index number associated with the session.
Shaping rate	Maximum bandwidth configured for the session. Each active queue on the session receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.

Sample Output

show class-of-service l2tp-session

```
user@host> show class-of-service l2tp-session 123
L2TP Session Username: user1@bng.com, Index: 12553
Physical interface: ge-2/0/0, Index: 130
```

```
Queues supported: 8, Queues in use: 4
  Scheduler map: GEN-SCHED-MAP-EF-65%, Index: 5212
  Shaping rate: 200000 bps
```


PART 4

Index

Index

R

S

T

U