JUNIPER NETWORKS

Junos® OS

CoS Bandwidth Feature Guide for Subscriber Services

Release
14.1

Published: 2014-04-25

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

Junos® OS CoS Bandwidth Feature Guide for Subscriber Services
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:

- 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 xv 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 xv 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**

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

Convention	Description	Examples
`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		
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**.

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

Convention	Description	Examples
> (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

Shaping-Rate Adjustments

Hierarchical CoS Shaping-Rate Adjustments Overview

This overview describes how MX Series 3D Universal Edge Routers installed in a subscriber access network can adjust hierarchical class-of-service (CoS) parameters to prevent bandwidth contention at subscriber interfaces.

Hierarchical CoS is supported only for subscriber interfaces on EQ DPC or MPC interfaces operating in hierarchical scheduler mode.

The characteristics of voice, data, and video applications vary widely in their requirements for traffic throughput, bandwidth management, delay and jitter tolerance, and buffer depth. To prevent bandwidth contention at subscriber interfaces, you can configure applications such as ANCP and Multicast to perform real-time adjustments to the shaping rate configured for subscriber interfaces for residential gateways. Enabling shaping-rate adjustments on the router can prevent bandwidth contention at the interface from causing degradation of the subscriber's voice, data, or video services.

Depending on the application, shaping-rate adjustments are supported on Enhanced Queueing (EQ) DPCs on MX Series routers and MPC/MIC modules on MX Series routers.

Types of Shaping-Rate Adjustments

The ANCP application supports *absolute* adjustments to a specific shaping-rate value. You can configure ANCP to communicate the subscriber local loop speed to the MX Series router, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. ANCP acquires subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS.

The OIF mapping and reverse OIF mapping multicast applications support *delta* adjustments that increase or decrease the current shaping rate by a certain value. The system adjusts traffic destined to the subscriber using reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

Levels of Shaping-Rate Adjustments

Both absolute and delta adjustments are made to a subscriber's aggregate shaping rate on a level 3 scheduler node.

Adjustments that occur on the scheduler node can also impact the shaping rates for all queues. This adjustment can be undesirable for service providers who want to provide a premium level of service on specific queues.

For delta-based adjustments by multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

Figure 1 on page 4 shows a sample multicast network with shaping rates adjusted at the scheduler node level. The shaping rate is reduced by 4 Mbps (from 41 Mbps to 37 Mbps) at the scheduler node for subscriber interface 1, which reduces the rates of both the best effort and video on demand (VoD) service queues.

Figure 1: Scheduler Node and Queues with Adjusted Shaping Rates

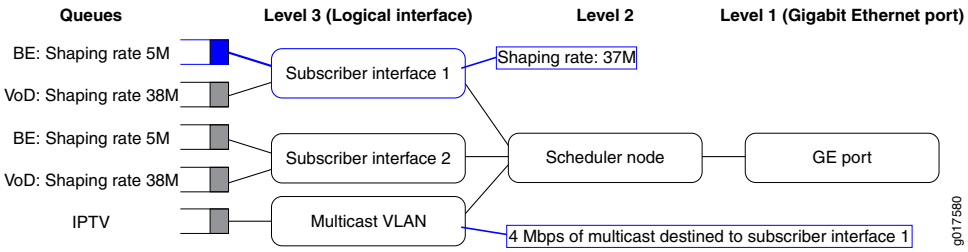

Figure 2 on page 4 shows the same network with queue-based adjustments enabled for the best-effort queue on subscriber 1. The shaping rate of the best-effort queue is reduced by 4 Mbps (from 5 Mbps to 1 Mbps). The VoD service queue is not affected.

Figure 2: Queue with Adjusted Shaping Rate

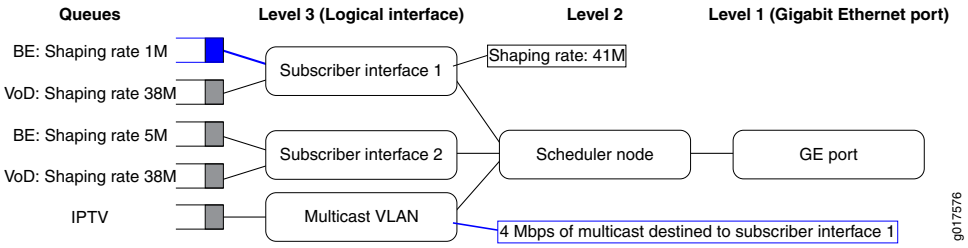

Related Documentation

- Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 37
- Configuring Shaping-Rate Adjustments on Queues on page 38
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 5
- Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping on page 45

- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 65

Shaping Rate Adjustments for Subscriber Local Loops Overview

This overview describes how an MX Series 3D Universal Edge Router installed as an edge router can adjust hierarchical CoS policy for subscriber interfaces for subscriber local loops. You can configure the router to throttle the traffic sent to subscriber local loops so that the traffic does not exceed the current data transmission rate of those lines. This feature ensures that changes to subscriber local loop speeds do not cause bandwidth contention at the subscriber's residential gateway.

In a typical subscriber access network, traffic destined to a subscriber is delivered from the access network, through an edge router, to a DSLAM. The DSLAM multiplexes subscriber traffic through a DSL, also known as a *local loop*, to the subscriber's residential gateway. When line noise or cross talk in a subcarrier causes the error rate on a DSL to exceed a certain threshold, the DSLAM can adapt itself by lowering the data transmission rate to that carrier device. A lower data transmission rate is less susceptible to induced errors.

You can configure an MX Series router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. Whenever a DSLAM resynchronizes a subscriber local loop speed, the router adjusts the configured shaping rate for that line so that the aggregate egress traffic to those subscribers is shaped to the local loop speed before the traffic reaches the DSLAM. Unless the maximum amount of bandwidth allocated to the subscriber interface on the router is throttled to the local loop speed, bandwidth contention can occur at the subscriber's residential gateway, which can cause the DSLAM to drop packets. This type of shaping-rate adjustment requires the topology discovery and traffic-monitoring features of the Access Node Control Protocol (ANCP).

You can configure ANCP to communicate the subscriber local loop speed to the MX Series router, which in turn throttles traffic destined to the associated subscriber interface so that it matches the subscriber local loop speed. ANCP acquires subscriber line rate information from DSLAMs and then communicates this data transmission rate for use with CoS.

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 33
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 40
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 45
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 65
- For more information about the ANCP protocol, see the *ANCP and the ANCP Agent Overview*.

CoS Adjustment Control Profiles Overview

CoS adjustment control profiles control which applications and algorithms can modify a subscriber's shaping characteristics after a subscriber is instantiated. Subscriber shaping characteristics are configured using the Junos OS CLI or by RADIUS messages. Adjustment control profiles enable subscriber shaping characteristics by to be adjusted by other applications like ANCP, PPPoE tags, and RADIUS Change of Authorization (CoA) after a subscriber is instantiated. Adjustment control profiles are router-wide and apply to both static and dynamic interfaces.

Table 3 on page 6 describes the terms used to define shaping characteristics.

Table 3: Terms and Definitions for Shaping Characteristics

Term	Definition
Shaping-rate	The maximum rate of a scheduler node or queue. Also known as Peak Information Rate (PIR).
Overhead-accounting mode	A class-of-service traffic-control profile attribute that specifies whether the downstream network is a frame-based network, like Ethernet, or a cell-based network, like ATM. NOTE: The downstream network is not necessarily the directly attached network. In typical broadband services router (BSR) network configurations, the directly attached network is an Ethernet access network, which provides access to either another frame-based network, or a cell-based network. When cell-mode is specified, the Juniper Networks router adjusts rates (like the shaping-rate) to "rate * 48/53" to account for 5-byte ATM AAL5 headers.
Overhead-accounting bytes	A class-of-service traffic-control profile attribute that specifies the number of bytes per packet to be included or excluded from the shaping mechanism. For example, to properly account for a 4-byte header stripped by the downstream network, set the overhead-accounting bytes to -4. To properly account for a 12-byte header added by the downstream network, set the overhead-accounting bytes to 12.
Effective-shaping-rate	The shaping-rate resulting from three attributes: shaping-rate, overhead-accounting mode, and overhead-accounting bytes.
Rate Adjustment	The adjustment that other applications like ANCP, PPPoE tagging, and RADIUS CoA can make when the shaping rate is configured through the CLI or RADIUS.

Effective Shaping Rate

CoS is responsible for communicating shaping rate information from the Routing Engine to the Packet Forwarding Engine. The shaping-rate, also known as peak information rate (PIR), is the maximum rate for a scheduler node or queue.

The true rate of a subscriber at the access-loop/CPE is a function of:

- The shaping-rate in effect for the subscriber's household, in bits per second.
- Whether the subscriber is connected to a frame-based or cell-based network. This attribute is known in CoS as the overhead-accounting mode.

- **overhead-accounting frame-mode-bytes**—When the **overhead-accounting** mode is set to **frame-mode**, this is the number of bytes in each frame that are accounted for by the shaper. This value represents the number of bytes that are encapsulated and decapsulated by the downstream equipment.
- **overhead-accounting cell-mode-bytes**—When overhead-accounting mode is set to **cell-mode**, this is the number of bytes in each frame that are accounted for by the shaper for a downstream cell-mode network. This value represents the number of bytes that are encapsulated and decapsulated by the downstream equipment.

You configure the values for the **shaping-rate** and **overhead-accounting** mode options under either the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level or the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level. These options are supported on MX Series routers. The applications and specified algorithms configured in the adjustment control profile use the values of these options to adjust the shaping rate for static and dynamically instantiated subscribers.

NOTE: Chassis *egress-shaping-overhead* is not included in the effective rate. Egress-shaping-overhead accounts for the physical interface overhead (ISO OSI Layer 1). Effective shaping-rate is a Layer 2 (ISO OSI) rate.

Applications and Associated Algorithms in Adjustment Control Profiles

Table 4 on page 7 describes the applications and their associated default algorithms that can be configured to perform rate adjustments after the subscriber is instantiated.

Table 4: Adjustment Control Profile Applications and Algorithms

Application	Default Priority	Default Algorithm	Description
RADIUS-CoA	1	Adjust-always	RADIUS CoA messages can update the subscriber's attributes (like shaping rate) after the subscriber is authenticated and QoS parameters (like shaping rate) are assigned.
ANCP	1	Adjust-always	The ANCP application can modify the existing shaping rate for both static and dynamic logical interfaces, and static interface sets. By default, ANCP can override all other applications. The shaping rate must be specified in order to override it.
PPPoE-Tags	2	Adjust-less	The PPPoE IA tag access-rate-downstream can modify the Junos OS CLI configured shaping- rate value, as well as the RADIUS shaping- rate value. By default, these values can be modified by subsequent RADIUS CoA messages and ANCP actions. These values are conveyed in PPPoE (PADI) discovery packets.

NOTE: The lower the priority value, the higher the priority.

You must enable each application to perform rate adjustments. Rate adjustments are global and affect all static and dynamically instantiated subscribers. The following rules apply to adjustment control profiles:

- If no adjustment control profile is configured, the default adjustment control profile is used.
- You can configure a maximum of one adjustment control profile; a commit error occurs if you configure more than one adjustment control profile.
- If an application is not configured with an adjustment control profile, Junos OS uses its default values for priority and algorithm. For example, if ANCP is not configured in the adjustment control profile, the ANCP application is set to a priority of 1 and the algorithm is set to adjust-always.
- Adjustment control profiles apply to both static and dynamic interfaces.
- You can configure the algorithm to the following values:
  - Adjust-never
  - Adjust-always
  - Adjust less
  - Adjust less than or equal
  - Adjust greater
  - Adjust greater than or equal
- When you modify an adjustment control profile, the changes take effect immediately and the modified profile is used for all further adjustments. However, existing adjustments are not reevaluated when you modify the adjustment control profile.

  For example, if you have an ANCP adjustment that overrides a PPPoE adjustment on interface ge-1/1/0.100, and then you use the adjustment control profile to change the priority so that the ANCP priority is now lower than the PPPoE priority, Junos OS does not go back and reevaluate the adjustment on ge-1/1/0.100.

Related Documentation

- Configuring CoS Adjustment Control Profiles on page 35
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- **adjustment-control-profiles on page 82**

CHAPTER 2

Traffic Distribution

Bandwidth Management for Downstream Traffic in Edge Networks Overview

In a subscriber access network, traffic with different encapsulations can be passed downstream to other customer premise equipment (CPE) through the MX Series router. Managing the bandwidth of downstream ATM traffic to Ethernet interfaces can be especially difficult because of the different Layer 2 encapsulations.

The *overhead accounting* feature enables you to shape traffic based on either frames or cells and assign a byte adjustment value to account for different encapsulations.

This feature is available on MIC and MPC interfaces.

Guidelines for Configuring the Shaping Mode

Frame shaping mode is useful for adjusting downstream traffic with different encapsulations. In frame shaping mode, shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead. Frame is the default shaping mode on the router.

Cell shaping mode is useful for adjusting downstream cell-based traffic. In cell shaping mode, shaping is based on the number of bytes in cells, and accounts for the cell encapsulation and padding overhead.

When you specify cell mode, the resulting traffic stream conforms to the policing rates configured in downstream ATM switches, reducing the number of packet drops in the Ethernet network.

To account for ATM segmentation, the MX Series router adjusts all of the rates by 48/53 to account for ATM AAL5 encapsulation. In addition, the router accounts for cell padding, and internally adjusts each frame by 8 bytes to account for the ATM trailer.

Guidelines for Configuring Byte Adjustments

When the downstream traffic has different byte sizes per encapsulation, it is useful to configure a *byte adjustment* value to adjust the frame sizes. For example, you can configure the frame shaping mode and a byte adjustment value to account for differences in Layer 2 protocols for downstream Ethernet traffic.

We recommend that you specify a byte adjustment value that represents the difference between the CPE protocol overhead and B-RAS protocol overhead.

The system rounds up the byte adjustment value to the nearest multiple of 4. For example, a value of 6 is rounded to 8, and a value of –10 is rounded to –8.

You do not need to configure a byte adjustment value to account for the downstream ATM network. However, you can specify the byte value to account for additional encapsulations or decapsulations in the downstream network.

Relationship with Other CoS Features

Enabling the overhead accounting feature affects the resulting shaping rates, guaranteed rate, and excess rate parameters, if they are configured.

The overhead accounting feature also affects the egress shaping overhead feature that you can configure at the chassis level. We recommend that you use the egress shaping-overhead feature to account for the Layer 2 overhead of the outgoing interface, and use the overhead-accounting feature to account for downstream traffic with different encapsulations and cell-based networks.

When both features are configured together, the total byte adjustment value is equal to the adjusted value of the overhead-accounting feature plus the value of the egress-shaping-overhead feature. For example, if the configured byte adjustment value is 40, and the router internally adjusts the size of each frame by 8, the adjusted overhead accounting value is 48. That value is added to the egress shaping overhead of 24 for a total byte adjustment value of 72.

Setting the Shaping-Rate and Overhead-Accounting Class-of-Service Attributes Based on Access Line Information

You can also use access line parameters in PPPoE discovery packets to set the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network. This feature is supported on MIC and MPC interfaces on MX Series routers.

The shaping rate is based on the actual-data-rate-downstream attribute.

The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).

You can mix ANCP and PPPoE vendor-specific tags for dynamically instantiated static and interface sets so that the shaping rate is first set using PPPoE vendor-specific tags and is later adjusted by ANCP. In this case, the shaping rate value overrides the PPPoE value.

The Access Node or DSLAM may forward access line information using several methods. This feature only uses access line information received in Vendor-Specific Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101].

When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured in the CLI for shaping-rate and overhead-accounting statements at the **[edit dynamic-profiles profile-name class-of-service traffic-control-profile]** hierarchy level. The shaping rate is based on the actual-data-rate-downstream attribute, and is only overridden if the vs-tag value is less than the configured value.

Related Documentation

- To configure overhead accounting for static Ethernet interfaces, see *Configuring Static Shaping Parameters to Account for Overhead in Downstream Traffic Rates*
- To configure overhead accounting for dynamic subscriber access, see Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 51
- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 17
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 46

Excess Bandwidth Distribution on MIC and MPC Interfaces Overview

Service providers often used tiered services to provide bandwidth for excess traffic as traffic patterns vary. By default, excess bandwidth between a configured guaranteed rate and shaping rate is shared equally among all queues on MIC and MPC interfaces, which might not be optimal for all subscribers to a service.

You can adjust this distribution by configuring the rates and priorities for the excess bandwidth.

By default, when traffic exceeds the shaping or guaranteed rates, the system demotes traffic with guaranteed high (GH) priority and guaranteed medium (GM) priority. You can disable this priority demotion for the MIC and MPC interfaces in your router.

Related Documentation

- *Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs*
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 52
- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- Traffic Burst Management on MIC and MPC Interfaces Overview on page 12

Traffic Burst Management on MIC and MPC Interfaces Overview

You can manage the impact of bursts of traffic on your network by configuring a burst-size value with the shaping rate or the guaranteed rate. The value is the maximum bytes of rate credit that can accrue for an idle queue or scheduler node. When a queue or node becomes active, the accrued rate credits enable the queue or node to catch up to the configured rate.

Figure 3: Sample Burst Shaping Rates

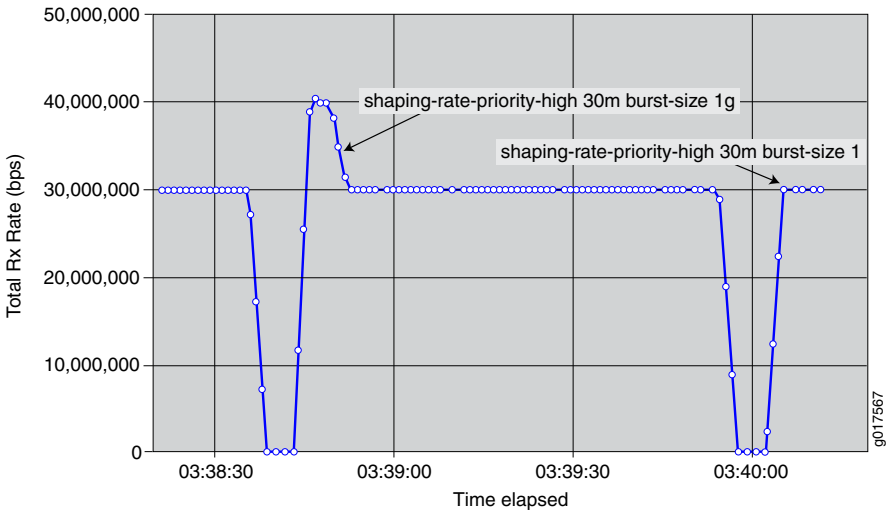

In Figure 3 on page 12, the network administrator configures a large burst-size value for the shaping rate, then configures a small burst-size value. The larger burst size is subject to a maximum value. The smaller burst size is subject to a minimum value that enables the system to achieve the configured rates.

In both configurations, the scheduler node can burst beyond its shaping rate for a brief interval. The burst of traffic beyond the shaping rate is more noticeable with the larger burst size than the smaller burst size.

- Guidelines for Configuring the Burst Size on page 12
- How the System Calculates the Burst Size on page 13

Guidelines for Configuring the Burst Size

Typically, the default burst-size (100 ms) for both scheduler nodes and queues on MIC and MPC interfaces is adequate for most networks. However, if you have intermediate equipment in your network that has very limited buffering and is intolerant of bursts of traffic, you might want to configure a lower value for the burst size.

Use caution when selecting a different burst size for your network. A burst size that is too high can overwhelm downstream networking equipment, causing dropped packets and inefficient network operation. Similarly, a burst size that is too low can prevent the network from achieving your configured rate.

When configuring a burst size, keep the following considerations in mind:

- The system uses an algorithm to determine the actual burst size that is implemented for a node or queue. For example, to reach a shaping rate of 8 Mbps, you must allocate 1Mb of rate credits every second. A shaping rate of 8 Mbps with a burst size of 500,000 bytes of rate-credit per seconds enables the system to transmit at most 500,000 bytes, or 4 Mbps. The system cannot implement a burst size that prevents the rate from being achieved.

  For more information, see "How the System Calculates the Burst Size" on page 13.

- There are minimum and maximum burst sizes for each platform, and different nodes and queue types have different scaling factors. For example, the system ensures the burst cannot be set lower than 1 Mbps for a shaping rate of 8 Mbps. To smoothly shape traffic, rate credits are sent much faster than once per second. The interval at which rate credits are sent varies depending on the platform, the type of rate, and the scheduler level.
- When you have configured adjustments for the shaping rate (either by percentage or through an application such as ANCP or Multicast OIF), the system bases the default and minimum burst-size calculations on the adjusted shaping rate.
- When you have configured cell shaping mode to account for ATM cell tax, the system bases the default and minimum burst-size calculations on the post-tax shaping rate.
- The guaranteed rate and shaping rate share the value specified for the burst size. If the guaranteed rate has a burst size specified, that burst size is used for the shaping rate; if the shaping rate has a burst size specified, that bursts size is used for the guaranteed rate. If you have specified a burst size for both rates, the system uses the lesser of the two values.
- The burst size configured for the guaranteed rate cannot exceed the burst-size configured for the shaping rate. The system generates a commit error.
- If you have not configured a guaranteed rate, logical interfaces and interface sets receive a default guaranteed rate from the port speed. Queues receive a default guaranteed rate from the parent logical interface or interface set.

How the System Calculates the Burst Size

When calculating the burst size, the system uses an exponent of a power of two. For example:

Shaping-rate in bps * 100 ms / (8 bits/byte * 1000 ms/s) = 1,875,000 bytes

The system then rounds this value up. For example, the system uses the following calculation to determine the burst size for a scheduler node with a shaping rate of 150 Mbps:

Max (Shaping rate, Guaranteed rate) bps * 100 ms / (8 bits/byte * 1000 ms/s) = 1,875,000 bytes

Rounded up to the next higher power of two = 2,097,150 (which is 221, or 0x200000)**

The system assigns a single burst size to each of the following rate pairs:

- Shaping rate and guaranteed rate
- Guaranteed high (GH) and guaranteed medium (GM)
- Excess high (EH) and excess low (EL)
- Guaranteed low (GL)

To calculate the burst size for each pair, the system:

- Uses the configured burst-size if only one of the pair is configured.
- Uses the lesser of the two burst sizes if both values are configured.
- Uses the next lower power of two.
- To calculate the minimum burst size, the system uses the greater of the two rates.

Related Documentation

- *Per-Priority Shaping on MIC and MPC Interfaces Overview*
- *Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs*

Distribution of Demux Subscribers in an Aggregated Ethernet Interface

This topic describes the distribution options available for demux subscriber interfaces over aggregated Ethernet.

Distribution Models

By default, the system supports hash-based distribution for all subscriber interface types in an aggregated Ethernet bundle configured without link protection. In this model, traffic for a logical interface can be distributed over multiple links in the bundle. This model is desirable when there are many flows through the logical interface and you need to load balance those flows.

Note that if the distribution flows are not even, egress CoS scheduling can be inaccurate. In addition, scheduler resources are required on every link of the aggregated Ethernet interface. For example, if subscriber traffic is allocated 10 MB for a triple-play service over four links in a bundle, each of the links could receive 2.5 MB of traffic. High-density services such as video could be limited by the bandwidth on one of the links.

Targeted distribution enables you to target the egress traffic for an IP or VLAN demux subscriber on a single member link, using a single scheduler resource. To achieve load balancing over the member links, the system distributes the subscriber interfaces equally among the links. This enables the subscriber that is allocated 10 MB to be accurately scheduled as the traffic flows through.

Sample Targeted Distribution Topology

Figure 4 on page 15 displays a sample targeted distribution of subscriber traffic across links in an aggregated Ethernet interface. A primary and backup link is allocated for each subscriber.

Figure 4: Targeted Subscriber Links

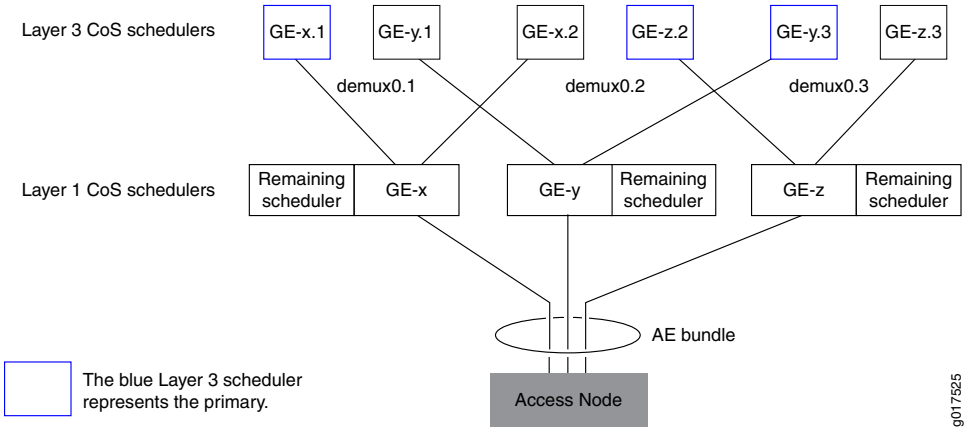

For example, if link **GE-x** went down, subscriber 1 can begin forwarding over the backup, which is link **Ge-y**. When link **GE-y** comes back up, subscriber 1 switches back to its primary link, **GE-x**.

In the event that both **GE-x** and **GE-y** go down, subscriber 3 starts forwarding through its backup, **GE-z**. Subscriber 1 will have lost its primary and backup links, and will also begin forwarding out the **GE-z** link. A new level 3 scheduler is assigned for this subscriber on link **GE-z**. If there is a momentary lapse between the time that a new scheduler is allocated and forwarding switches to **GE-z**, the traffic will be forwarding through to the remaining scheduler. Subscriber 2 continues to forward through its primarily link, **GE-z**.

Redundancy and Redistribution Mechanisms

Two types of redundancy are available in the targeted distribution model: link redundancy and module redundancy.

By default, an aggregated Ethernet interface is enabled with link redundancy. Backup links for a subscriber are chosen based on the link with the least number of subscribers, which provides redundancy if a link fails.

The module redundancy option enables you to provide redundancy if a module or a link fails. Backup links for a subscriber are chosen on a different DPC or MPC from the primary link, based on the link with the least number of subscribers among the links on different modules. You can enable this for the aggregated Ethernet interface.

When links are removed, affected subscribers are redistributed among the active remaining backup links. When links are added to the system, no automatic redistribution occurs. New subscribers are assigned to the links with the fewest subscribers (which are typically the new links).

Considerations and Best Practices

Keep the following guidelines in mind when configuring targeted distribution for demux subscribers:

- You can manage subscribers with both hash-based and targeted distribution models in the same network. For example, you can allocate subscribers with interface types such as PPPoE with hash-based distribution, and enable demux subscribers with targeted distribution.
- We recommend that you configure module redundancy to protect against module failures. When module redundancy is enabled, you can ensure an even distribution of subscribers if you allocate no more than 50 percent of the links on a single DPC or MPC.
- During normal network operations, the system maintains an even balance of subscribers among the links in a bundle, even as subscribers log in and out. However, if the distribution of a bundle becomes uneven (for example, when a link goes down and new subscribers are logging in), you can perform a manual rebalance of the bundle. In addition, you can configure periodic rebalancing of the bundle with a specific time interval.
- When you anticipate that a link will be down for an extended time, and you want to ensure that backup links are provisioned for all subscribers, we recommend that you remove the failed link from the bundle. This forces the affected subscribers to redistribute to other links.
- We recommend that you apply a remaining traffic-control profile to the logical interface to ensure that minimal scheduling parameters are applied to the remaining subscriber traffic. This provides scheduling for subscribers that do not have schedulers allocated because they have not been configured or they have been over-provisioned, or because of scheduler transitions on multiple link failures.
- If you perform a cold restart on the router when it is forwarding active subscribers, the subscriber interfaces with targeted distribution are assigned to the first links that become available when the system is initializing so forwarding can begin. To rebalance the system following a cold restart, perform a manual rebalance of the bundle. In addition, we recommend that you configure Graceful Routing Engine switchover (GRES) on the router to enable nonstop forwarding during switchover, and avoid performing cold restarts.
- To ensure appropriate and predictable targeted distribution, you must configure chassis network services to use **enhanced-ip** mode.
- Unless specifically separated, multicast traffic egresses in parallel with unicast traffic, sharing the CoS hierarchy and aggregated Ethernet flow distribution.

Related Documentation

- *Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces*
- *Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface*
- *Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface*
- *Static or Dynamic Demux Subscriber Interfaces over Aggregated Ethernet Overview*

Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags

You can use access line parameters in PPPoE discovery packets to set the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network. This feature is supported on MPC/MIC interfaces on MX Series routers.

The shaping rate is based on the actual-data-rate-downstream attribute.

The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).

You can configure class-of-service attributes, for example the shaping-rate, using the CLI, RADIUS vendor-specific attributes, ANCP, multicast, or in this case, PPPoE vendor-specific tags.

CLI Interaction with PPPoE Vendor-Specific Tags

When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured in the CLI for the **shaping-rate** and **overhead-accounting** statements at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level. The shaping rate is based on the actual-data-rate-downstream attribute, and is only overridden if the vs-tag value is less than the configured value.

To enable this feature, include the **actual-data-rate-downstream** or **access-loop-encapsulation** option with the **vendor-specific-tags** statement at the **[edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]** hierarchy level.

RADIUS Interaction with PPPoE Vendor-Specific Tags

When you enable this feature, the PPPoE vendor-specific tags override the dynamic configuration of the shaping-rate and overhead-accounting values in RADIUS vendor-specific attributes. The shaping-rate value is only overridden if the vs-tag value is less than the RADIUS value.

RADIUS CoA can overwrite the existing values. Upon receipt of a RADIUS CoA, the RADIUS value overrides the value set from the PPPoE vendor-specific tags.

PPPoE vendor-specific tags can override the RADIUS values, but a later RADIUS CoA request can then override that value.

ANCP Interaction with PPPoE Vendor-Specific Tags

You can mix ANCP and PPPoE vendor-specific tags on dynamic PPPoE interfaces, dynamically instantiated PPPoE interfaces, and ACI-sets. ANCP values override the PPPoE values. In this case, the ANCP shaping rate value overrides the PPPoE value.

Multicast QoS Adjustment Interaction with PPPoE Vendor-Specific Tags

Multicast QoS adjustments are not affected by this feature. The multicast adjustments adjust the shaping-rate set by PPPoE vendor-specific tags.

Shaping Rate Restrictions

Shaping rate has the following restrictions regarding the downstream-rate:

- If the downstream-rate is less than the configured shaping-rate (as set in the CLI or using RADIUS attributes) then it is applied, subject to other restrictions. If the downstream-rate is greater than or equal to the configured shaping-rate, no changes are performed.
- The downstream-rate cannot be less than a configured guaranteed-rate. If it is, the downstream-rate is set to the guaranteed-rate.
- The downstream-rate cannot be less than a configured adjust-minimum-rate. If it is, the downstream-rate is set to the adjust-mimimum-rate.
- The downstream-rate cannot be less than 1000 bps. If it is, the downstream-rate is set to 1000 bps.
- The downstream-rate cannot be less than the sum of the transmit-rates of all queues.

Related Documentation

- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 46

CHAPTER 3

Scaling

Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview

The 30-Gigabit Ethernet Queuing and 60-Gigabit Ethernet Queuing and Enhanced Queuing Ethernet Modular Port Concentrators (MPCs) provide a set of dedicated queues for subscriber interfaces configured with hierarchical scheduling or per-unit scheduling.

The dedicated queues offered on these MPCs enable service providers to reduce costs through different scaling configurations. For example, the 60-Gigabit Ethernet Enhanced Queuing MPC enables service providers to reduce the cost per subscriber by allowing many subscriber interfaces to be created with four or eight queues. Alternatively, the 30-Gigabit Ethernet and 60-Gigabit Ethernet Queuing MPCs enable service providers to reduce hardware costs, but allow fewer subscriber interfaces to be created with four or eight queues.

This topic describes the overall queue, scheduler node, and logical interface scaling for subscriber interfaces created on these MIC and MPC combinations.

Queue Scaling for MIC and MPC Combinations

Table 5 on page 19 lists the number of dedicated queues and number of subscribers supported per MPC.

Table 5: Dedicated Queues for MIC and MPC Interfaces

MPC	Dedicated Egress Queues	Supported Subscriber Interfaces	Logical Interfaces with 4 Queues	Logical Interfaces with 8 Queues
30-Gigabit Ethernet Queuing MPC	64,000	16,000	16,000 (8000 per PIC)	8000 (4000 per PIC)
60-Gigabit Ethernet Queuing MPC	128,000	32,000	32,000 (8000 per PIC)	16,000 (4000 per PIC)

Table 5: Dedicated Queues for MIC and MPC Interfaces *(continued)*

MPC	Dedicated Egress Queues	Supported Subscriber Interfaces	Logical Interfaces with 4 Queues	Logical Interfaces with 8 Queues
60-Gigabit Ethernet Enhanced Queuing MPC	512,000	64,000	64,000 (16,000 per PIC)	64,000 (16,000 per PIC)

MPCs vary in the number of Packet Forwarding Engines on board. MPC1s, such as the 30-Gigabit Ethernet MPC, have one Packet Forwarding Engine. MPC2s, such as the 60-Gigabit Ethernet MPC, have two Packet Forwarding Engines. Each Packet Forwarding Engine has two schedulers that share the management of the queues.

A scheduler maps to one-half of a MIC; in CLI configuration statements, that one-half of a MIC corresponds to PIC 0, 1, 2, or 3. MIC ports are partitioned equally across the PICs. A two-port MIC has one port per PIC. A four-port MIC has two ports per PIC.

Each interface-set uses eight queues from total available egress queues.

Distribution of Queues on 30-Gigabit Ethernet Queuing MPCs

On 30-Gigabit Ethernet Queuing MPCs, each scheduler maps to different PICs. When only one MIC is installed, scheduler 0 maps to PIC 0 and scheduler 1 maps to PIC 1 on the MIC. When two MICs are installed, scheduler 0 can additionally distribute queues to PIC 2 on MIC 1, and scheduler 1 can additionally distribute queues to PIC 3 on MIC 1. However, the distribution of queues to the MICs is not hard-partitioned for 30-Gigabit Ethernet Queuing MPCs or other MPC1s. Distribution depends instead on how you allocate the queues to the PICs.

Figure 5 on page 20 shows the queue distribution on a 30-Gigabit Ethernet Queuing MPC with only one MIC installed. All 64,000 egress queues on the MPC are available to the single Packet Forwarding Engine. On the Packet Forwarding Engine, half of these queues (32,000) are managed by each scheduler. Scheduler 0 contributes all of its 32,000 queues to PIC 0. Scheduler 1 contributes all of its 32,000 queues to PIC 1.

Figure 5: Distribution of Queues on the 30-Gigabit Ethernet Queuing MPC with One MIC

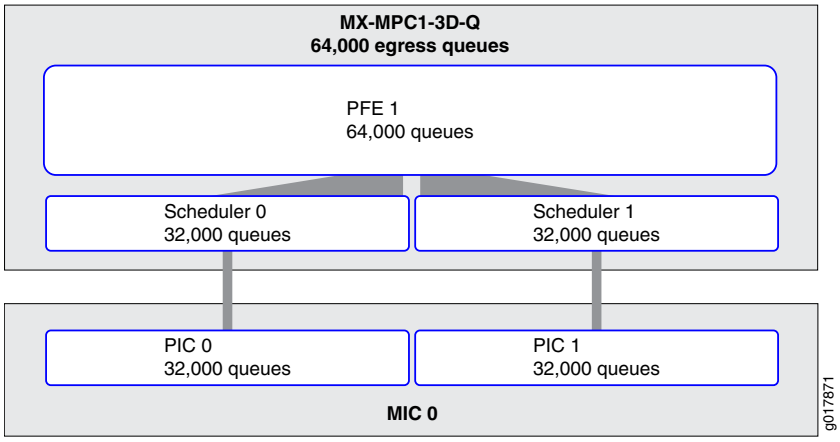

Figure 6 on page 21 shows the queue distribution on the same MPC with two MICs installed. In this case, each scheduler can supply two PICS, one on each MIC. Because the distribution of the queues across the MICs is not hard-partitioned, you can allocate from 0 to 32,000 queues from each scheduler's pool across the scheduler's associated PICs. For example, you can allocate 32,000 queues from Scheduler 0 to PIC 0, 4000 queues from Scheduler 1 to PIC 1, and 28,000 queues from Scheduler 1 to PIC 3. Alternatively, you can allocate the queues evenly across the PICs, or allocate them in other combinations with the limitation of 32,000 queues per PIC and 32,000 queues per port.

Figure 6: Distribution of Queues on the 30-Gigabit Ethernet Queuing MPC with Two MICs

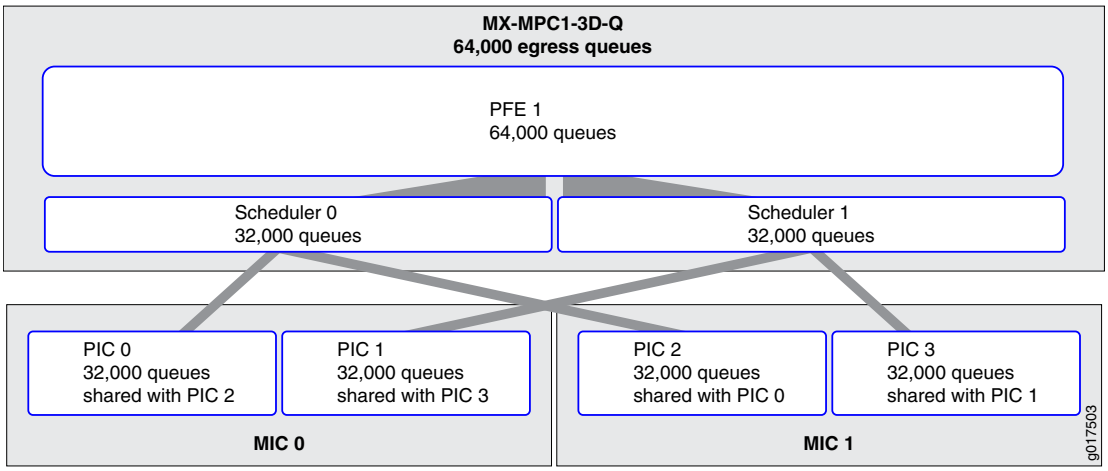

Distribution of Queues on 60-Gigabit Ethernet MPCs

On 60-Gigabit Ethernet Queuing and Enhanced Queuing Ethernet MPCs, each scheduler maps to a single PIC: PIC 0 or PIC 1 on MIC 0 and PIC 2 or PIC 3 on MIC 1. The distribution of the queues is hard-partitioned for these MPCs and other MPC2s; the only difference in distribution is in the total number of queues available.

For example, Figure 7 on page 22 shows how queues are distributed on a 60-Gigabit Ethernet Enhanced Queuing MPC. Of the 512,000 egress queues on the MPC, half (256,000) are available to each of the two Packet Forwarding Engines. On each Packet Forwarding Engine, half of these queues (128,000) are managed by each scheduler. The complete scheduler complement (128,000) is available to only one PIC in a MIC. Thus the total number of queues available depends on the number of MICs installed. The MPC must have 2 MICs to achieve the maximum of 512,000 queues. With a single MIC, the MPC can achieve only 256,000 queues.

Figure 7: Distribution of Queues on the 60-Gigabit Ethernet Enhanced Queuing MPC

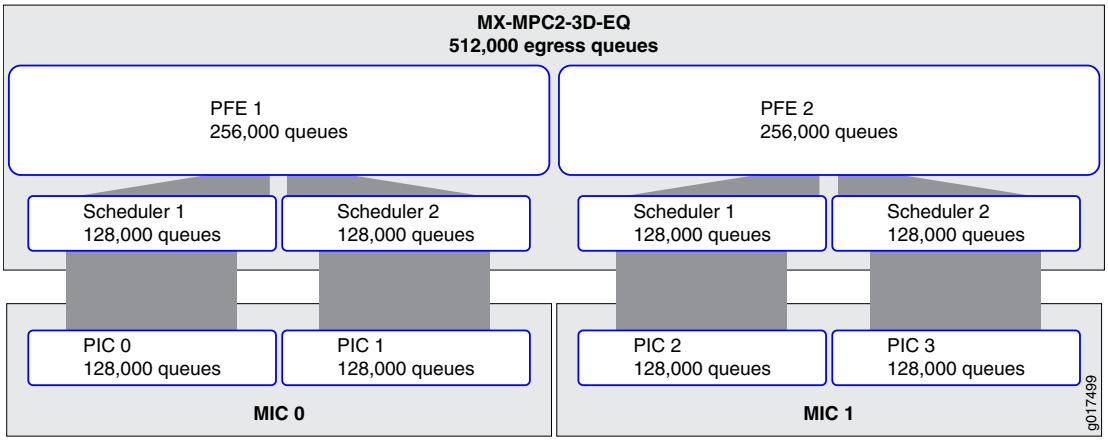

Determining Maximum Egress Queues and Subscriber Interfaces per Port

The number of MICs installed in an MPC and the number of ports per MIC do not affect the maximum number of queues available on a given port. These factors affect only how you are able to allocate queues (and, therefore, subscribers) for your network.

For example, a 30-Gigabit Ethernet Queuing MPC supports a maximum of 16,000 subscriber interfaces and has a maximum of 32,000 queues available per PIC. On this card, you can allocate up to 32,000 queues to a single port in each PIC. If you dedicate 4 queues per subscriber interface, you can accommodate a maximum of 8000 subscriber interfaces on a single port, and therefore need at least two ports to reach the maximum 16,000 subscriber interfaces. If you dedicate 8 queues per subscriber interface, you can accommodate a maximum of 4000 subscriber interfaces on a single port, and you need 4 ports for the maximum of 16,000 subscriber interfaces.

The 60-Gigabit Ethernet Enhanced Queuing MPC supports a maximum of 64,000 subscriber interfaces and has a maximum of 128,000 queues per PIC. You can allocate up to 128,000 queues to a single port in each PIC. However, if you dedicate 4 queues per subscriber interface, you can accommodate a maximum of only 16,000 subscriber interfaces on a single MPC port—not 32,000—because the 60-Gigabit Ethernet Enhanced Queuing MPC is limited to 16,000 subscriber interfaces per PIC. If you dedicate 8 queues per subscriber interface, you can also accommodate a maximum of 16,000 subscriber interfaces on a single MPC port. In either case, you need at least 4 ports to reach the maximum of 64,000 subscriber interfaces.

Managing Remaining Queues

When the number of available dedicated queues on the MPC drops below 10 percent, an SNMP trap is generated to notify you .

When the maximum number of dedicated queues on the MPCs is reached, a system log message, **COSD_OUT_OF_DEDICATED_QUEUES**, is generated. The system does not provide subsequent subscriber interfaces with a dedicated set of queues. For per-unit scheduling configurations, there are no configurable queues remaining on the MPC.

For hierarchical scheduling configurations, remaining queues are available when the maximum number of dedicated queues is reached on the MPC. Traffic from these logical interfaces are considered unclassified and attached to a common set of queues that are shared by all subsequent logical interfaces. These common queues are the default port queues that are created for every port. You can configure a traffic-control profile and attach that to the interface to provide CoS parameters for the remaining queues.

For example, when the 30-Gigabit Ethernet Queuing MPC is configured with 32,000 subscriber interfaces with four queues per subscriber, the MPC can support 16,000 subscribers with a dedicated set of queues. You can provide CoS shaping and scheduling parameters to the remaining queues for those subscriber interfaces by attaching a special traffic-control profile to the interface.

These subscriber interfaces remain with this traffic-control profile, even if dedicated queues become available.

Related Documentation

- For information about managing dedicated queues in a static CoS configuration, see *Managing Dedicated and Remaining Queues for Static CoS Configurations on MIC and MPC Interfaces*
- For information about managing dedicated queues in a dynamic subscriber access configuration, see Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces on page 57
- *Understanding Hierarchical Scheduling for MIC and MPC Interfaces*
- COSD System Log Messages on page ?

PART 2

Configuration

CHAPTER 4

Best Practices

Guidelines for Configuring Dynamic CoS for Subscriber Access

This topic describes the hardware requirements and guidelines for configuring dynamic CoS in a subscriber access environment.

Hardware Requirements for Dynamic CoS

Table 6 on page 28 lists the hardware requirements based on subscriber interface type for the hierarchical scheduling and per-unit scheduling dynamic CoS configurations.

Table 6: Hardware Required for Dynamic CoS Configurations

Dynamic CoS Configuration	Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC/MIC Modules on MX Series Routers	IQ2 PICs on M120 and M320 Routers	IQ2E PICs on M120 and M320 Routers
Hierarchical CoS	Static and dynamic VLANs	Yes	Yes	No	No
	Static and dynamic VLANs over aggregated Ethernet	Yes	Yes	No	No
	Static or dynamic IP demux interfaces	Yes	Yes	No	No
	Static or dynamic IP demux interfaces over aggregated Ethernet	Yes	Yes	No	No
	Static or dynamic VLAN demux interfaces	No	Yes	No	No
	Static or dynamic VLAN demux interfaces over aggregated Ethernet	No	Yes	No	No
	Static PPPoE interfaces	No	Yes	Yes	Yes
	Dynamic PPPoE interfaces	No	Yes	No	Yes
	Static or dynamic PPPoE interfaces over aggregated Ethernet	No	Yes	No	No
	L2TP LAC tunnel over PPP	No	Yes	No	No
	L2TP LNS inline service over PPP	No	Yes	No	No

Table 6: Hardware Required for Dynamic CoS Configurations *(continued)*

Dynamic CoS Configuration	Subscriber Interface Type	EQ DPCs on MX Series Routers	MPC/MIC Modules on MX Series Routers	IQ2 PICs on M120 and M320 Routers	IQ2E PICs on M120 and M320 Routers
Per-unit scheduling	Static and dynamic VLANs	Yes	Yes	No	No
	Static and dynamic VLANs over aggregated Ethernet	No	No	No	No
	Static or dynamic IP demux interfaces	Yes	No	No	No
	Static or dynamic IP demux interfaces over aggregated Ethernet	No	No	No	No
	Static or dynamic VLAN demux interfaces	No	No	No	No
	Static or dynamic VLAN demux interfaces over aggregated Ethernet	No	No	No	No
	Static PPPoE interfaces	No	Yes	Yes	Yes
	Dynamic PPPoE interfaces	No	No	Yes	Yes
	Static or dynamic PPPoE interfaces over aggregated Ethernet	No	No	No	No
	L2TP LAC tunnel over PPP	No	No	No	No
	L2TP LNS inline service over PPP	No	No	No	No

Configuration Guidelines for Dynamic Scheduling and Queuing

When configuring scheduling and queuing for subscriber access, consider the following guidelines:

- To improve CoS performance in IPv4, IPv6, and dual-stack networks that use a DHCP access model, we recommend that you use the **aggregate-clients replace** statement rather than the **aggregate-clients merge** statement.
- You can configure dynamic CoS with one of the following scheduling configurations:
  - For hierarchical scheduling configurations, you must enable hierarchical scheduling in the static CLI for the interface referenced in the dynamic profile. If not, the dynamic profile fails.
  - For per-unit scheduling configurations, you must enable per-unit scheduling in the static CLI for the interface referenced in the dynamic profile. If not, the dynamic profile fails and schedulers are not attached to the interface.
- You configure the traffic scheduling and shaping parameters in a traffic-control profile within the dynamic profile. You can configure the scheduler map and schedulers in a dynamic profile or in the **[edit class-of-service]** hierarchy. You must statically configure the remaining CoS parameters, such as hierarchical scheduling, classifiers, drop profiles, and forwarding classes, in the **[edit class-of-service]** hierarchy.
- You can configure only one traffic-control-profile under a dynamic profile.
- You must define the output-traffic-control-profile that binds the traffic-control profile to the interface within the same dynamic profile as the interface.
- We recommend that you provide different names for the schedulers defined in dynamic profiles that are used for access and services. For example, if there are two dynamic profiles, voice-profile and video-profile, provide unique names for the schedulers defined under those profiles.
- You must use a service dynamic profile with a different profile name for each RADIUS CoA request over the same logical interface.
- When you configure scheduler and scheduler map sharing in client profiles, schedulers and scheduler maps must use the unique ID format. If the client profile uses the unique ID format and you want to have either scheduler or scheduler map sharing for service activation, you must configure the service profile in unique ID format.

Configuration Guidelines for Dynamic Classifiers and Rewrite Rules

When you configure classifiers and rewrite rules for subscriber access, consider the following guidelines:

- To apply classifiers and rewrite rules to a subscriber interface in a dynamic profile, you must configure the rewrite rule and classifier definitions in the static **[edit class-of-service]** hierarchy and reference them in the dynamic profile.

- If a static classifier or a rewrite rule definition that is referenced by a dynamic subscriber interface does not exist, the configuration is invalid and the subscriber cannot log in.
- If a network administrator changes the static classifiers and rewrite rules definitions that are referenced in a dynamic profile with an active subscriber interface logged in, the changes are applied to the active subscriber interface immediately.
- If a network administrator deletes a classifier or a rewrite rule definition that is referenced by an active dynamic subscriber interface, the system removes the classifier or rewrite rule binding from the interface. The classifier is replaced by the default classifier. If the network administrator adds the removed classifier or rewrite rule to the configuration while the dynamic interface is active, the addition does not take effect until the subscriber logs out and then logs in again.

- IP demux interfaces can only instantiate Layer 3 rules (both rewrite rules and classifiers).
  - An IP demux subscriber interface can implicitly inherit a classifier from the underlying interface. If an IP demux interface is created without a classifier and a Layer 2 classifier is attached to the underlying interface, the IP demux interface also inherits the Layer 2 classifier. The show class-of-service interface ***interface-name*** command does not display this attachment.

    Table 7 on page 31 lists the classification rule configuration for an IP demux subscriber interface with a VLAN underlying interface.

Table 7: IP Demux Classification Rules

VLAN Underlying Interface Classifier Configuration	IP Demux Interface Classifier Configuration	Resulting Classifier Configuration
Layer 2	—	VLAN Layer 2
Layer 2	Layer 3	Demux Layer 3
Layer 3	—	Default
Layer 3	Layer 3	Demux Layer 3

  - An IP demux subscriber interface explicitly inherits Layer 2 rewrite rules from the underlying interface if a Layer 2 rewrite rule is present. The show class-of-service interface ***interface-name*** command displays the attachment.

    Table 8 on page 31 lists the rewrite rule configuration for an IP demux subscriber interface with a VLAN underlying interface.

Table 8: IP Demux Rewrite Rules

VLAN Underlying Interface Rewrite Rule Configuration	IP Demux Interface Rewrite Rule Configuration	Resulting Rewrite Rule Configuration
Layer 2	—	VLAN Layer 2

Table 8: IP Demux Rewrite Rules *(continued)*

VLAN Underlying Interface Rewrite Rule Configuration	IP Demux Interface Rewrite Rule Configuration	Resulting Rewrite Rule Configuration
Layer 2	Layer 3	VLAN Layer 2 and demux Layer 3
Layer 3	—	Default
Layer 3	Layer 3	Demux Layer 3

- An L2TP subscriber interface can implicitly inherit a classifier from the underlying interface.

  Table 9 on page 32 lists the classification rule configuration for an L2TP LAC subscriber interface with a VLAN underlying interface.

Table 9: L2TP Classification Rules

VLAN Underlying Interface Classifier Configuration	L2TP LAC Classifier Configuration	Resulting Classifier Configuration
Layer 2 or Fixed	Layer 2 or Fixed	VLAN Layer 2 or Fixed
Layer 2 or Fixed	Layer 3	Demux/PPPoE Layer 3
Layer 3	Layer 2 or Fixed	VLAN Layer 2 or Fixed
Layer 3	Layer 3	Demux/PPPoE Layer 3

- An L2TP LAC subscriber interface explicitly inherits Layer 2 rewrite rules from the underlying interface if a Layer 2 rewrite rule is present. Table 10 on page 32 lists the rewrite rule configuration for an L2TP LAC subscriber interface with a VLAN underlying interface.

Table 10: L2TP LAC Rewrite Rules

VLAN Underlying Interface Rewrite Rule Configuration	L2TP Interface Rewrite Rule Configuration	Resulting Rewrite Rule Configuration
Layer 2	Layer 2	VLAN Layer 2
Layer 2	Layer 3	VLAN Layer 2 and demux/PPPoE Layer 3
Layer 3	Layer 2	VLAN Layer 2 and demux/PPPoE Layer 3
Layer 3	Layer 3	Demux/PPPoE Layer 3

Configuration Guidelines for Dynamic CoS on Specific Interface Types

To obtain configuration guidelines for CoS on specific interface types, see:

- *CoS for Aggregated Ethernet Subscriber Interfaces Overview*
- *CoS for L2TP LAC Subscriber Interfaces Overview*
- *CoS for L2TP LNS Inline Services Overview*
- *CoS for PPPoE Subscriber Interfaces Overview*
- *CoS for Interface Sets of Subscribers Overview*

Related Documentation

- *CoS for Subscriber Access Overview*
- *Understanding Two-Level and Three-level Hierarchical CoS for Subscriber Interfaces on MX Series Routers*
- *Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Configuring Per-Unit Scheduling in a Dynamic Profile for Subscriber Access*
- *Configuring Static CoS for an L2TP LNS Inline Service*

Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops

These guidelines apply to configuring an MX Series 3D Universal Edge Router installed as an edge router to adjust the configured shaping rates on scheduler nodes for subscriber interfaces that represent subscriber local loops. This shaping-rate feature uses the topology discovery and traffic-monitoring features of ANCP.

When you enhance hierarchical CoS policy by configuring ANCP-driven shaping-rate adjustments, consider the following guidelines:

- Shaping-rate adjustments are supported on EQ DPCs and MPCs on MX Series routers.
- Shaping-rate adjustments are supported only for subscriber local loops that terminate at DSLAMs that you have configured as ANCP neighbors of the MX Series router.
- Shaping-rate adjustments are supported only for scheduler nodes for which you have configured an initial shaping rate by including the **shaping-rate** statement in a traffic-control profile applied to the scheduler node. Specify the initial shaping rate as a peak rate, in bits per second (bps), and not as a percentage. Other methods of configuring a shaping rate are not supported with this feature.
- Shaping-rate adjustments are supported only for scheduler nodes that are static logical interface sets that you have configured to operate at Level 3 of the scheduler hierarchy on the router. If an interface set is configured with a logical interface (such as unit 0) and queue, then the interface set is an internal scheduler node (as opposed to a root node or a leaf node) at Level 2 of the hierarchy. However, if there are no traffic-control

profiles are configured on logical interfaces in an interface set, then the interface set is an internal scheduler node at Level 3 of the hierarchy.

- Shaping-rate adjustments are supported only for subscriber interfaces over physical interfaces that you have configured to operate in hierarchical scheduler mode. Only ports on EQ DPCs in MX Series routers support hierarchical scheduler mode.
- After shaping-rate adjustments are enabled and the router has performed shaping-rate adjustments on a scheduler node, you can configure a new shaping rate by including the **shaping-rate** statement in a traffic-control profile and then applying that profile to that scheduler node. However, this new shaping-rate value does not immediately result in shaping traffic at the new rate. The scheduler node continues to be shaped at rate set by ANCP. Only when the ANCP shaping-rate adjustment feature is disabled is the scheduler node shaped at the newly configured shaping-rate.
- The Layer 2 Tunneling Protocol (L2TP) is often used to carry traffic securely between an L2TP Network Server (LNS) and an L2TP Access Concentrator (LAC). The QoS adjustment feature supports the shaping overhead options that you can use to add a specified number of bytes to the actual packet length when determining shaped session packet length. ANCP shaping-rate adjustments are not supported for ingress traffic, only for egress traffic. To configure the number of bytes to add to the packet at the egress side of the tunnel, include the **egress-shaping-overhead** and **mode** statements at the **[edit chassis fpc *slot-number* pic *pic-number* traffic-manager]** hierarchy level. Use the shaping overhead options if you need to account for encapsulation overhead.

For more information about the ANCP protocol, see the *ANCP and the ANCP Agent Overview*.

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 5
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 40
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 45
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 65

CHAPTER 5

Configuration Tasks for Adjustment Control Profiles

Configuring CoS Adjustment Control Profiles

To configure adjustment control profiles:

NOTE: You can only configure one adjustment control profile.

1. Configure the adjustment control profile name.

   [edit]
   user@host# **edit class-of-service adjustment-control-profiles** ***profile-name***

2. (Optional) Configure the adjustment controls for the Access Node Control Protocol (ANCP) application:

   [edit class-of-service adjustment-control-profiles *profile-name*]
   user@host# **set application ancp priority** ***priority*** **algorithm** ***algorithm***

3. (Optional) Configure the adjustment controls for the RADIUS CoA application:

   [edit class-of-service adjustment-control-profiles *profile-name*]
   user@host# **set application radius-coa priority** ***priority*** **algorithm** ***algorithm***

4. (Optional) Configure the adjustment controls for the PPPoE tags:

   [edit class-of-service adjustment-control-profiles *profile-name*]
   user@host# **set application pppoe-tags priority** ***priority*** **algorithm** ***algorithm***

5. (Optional) Verify your configuration.

```
user@host> show class-of-service adjustment-control-profiles
name: ANCP, priority: 1, algorithm: less;
name: RADIUS CoA, priority: 1, algorithm: always;
name: PPPoE IA tags, priority: 2, algorithm: less;
```

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- **adjustment-control-profiles on page 82**

- overhead-accounting (Dynamic Traffic Shaping) on page 103

CHAPTER 6

Configuration Tasks for Shaping-Rate Adjustments

Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers

Overview

Absolute adjustments and delta adjustments are performed at the scheduler node level. You can configure a minimum adjusted shaping rate at the scheduler node level using static or dynamic CoS parameters.

This feature is supported for adjustments performed by the ANCP and multicast applications on both EQ DPCs and MPC/MIC modules on MX Series routers.

BEST PRACTICE: For multicast traffic, you can configure a minimum adjusted shaping rate at the queue level. We recommend that you configure the minimum adjusted value at the scheduler node or the queue, but not both.

When you configure a minimum adjusted value for a node and for a scheduler that is referenced by a scheduler map in the same traffic-control-profile, the system uses the minimum value from the scheduler.

This feature is supported for adjustments performed by the ANCP and multicast applications on both EQ DPCs and MPC/MIC modules on MX Series routers.

Configuring a Static Minimum Adjusted Shaping Rate on Scheduler Nodes

To apply a minimum adjusted shaping rate for a scheduler node:

- Configure the **adjust-minimum** statement for the static traffic-control profile.

  [edit class-of-service traffic-control-profiles *profile-name*]
  user@host# **set adjust-minimum** *rate*

Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes

To apply a minimum adjusted shaping rate for a scheduler node:

- Configure the **adjust-minimum** statement for the dynamic traffic-control profile.

  [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]
  user@host# **set** adjust-minimum *rate*

Related Documentation

- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 117
- Configuring Shaping-Rate Adjustments on Queues on page 38
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3

Configuring Shaping-Rate Adjustments on Queues

Overview

By default, the multicast application adjusts the shaping rates at the scheduler node level. This adjustment also impacts the shaping rates for all queues, which can be undesirable for service providers who want to provide a premium level of service on specific queues.

For multicast applications, you can control the distribution of shaping rates among queues by assigning the percentage of adjustment allowed for each queue. In addition, you can set a minimum adjusted shaping rate for each queue.

This feature is supported for adjustments performed by the multicast application on MPC/MIC modules on MX Series routers.

BEST PRACTICE: We recommend that you configure the minimum adjusted value at the scheduler node or the queue, but not both.

When you configure a minimum adjusted value for a node and for a scheduler that is referenced by a scheduler map in the same traffic-control-profile, the system uses the minimum value from the scheduler.

This feature is supported for adjustments performed by the multicast application on MPC/MIC modules on MX Series routers.

Configuring a Static Shaping-Rate Adjustment for Queues

To configure adjustment parameters for a queue:

1. Configure the percentage of adjustment for the shaping rate.

   [edit class-of-service schedulers *scheduler-name*]
   user@host# **set adjust-percent *percentage***

2. Configure the minimum adjusted value for the shaping rate.

   Do one of the following:

   - Configure the minimum adjusted value for the queue.

     [edit class-of-service schedulers *scheduler-name*]
     user@host# **set adjust-minimum *rate***

   - Configure the minimum adjusted value for the node.

     [edit class-of-service traffic-control-profile *profile-name*]
     user@host# **set adjust-minimum *rate***

BEST PRACTICE: Ensure that the minimum adjusted value that you configure does not exceed the shaping rate and is not lower than the configured transmit rate.

Configuring a Dynamic Shaping-Rate Adjustment for Queues

To configure adjustment parameters for a queue in a dynamic profile:

1. Configure the percentage of adjustment for the shaping rate.

   [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
   user@host# **set** adjust-percent ***percentage***

2. Configure the minimum adjusted value for the shaping rate.

   Do one of the following:

   - Configure the minimum adjusted value for the queue.

     [edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]
     user@host# **set** adjust-minimum **(*rate* | $junos-cos-adjust-minimum)**

- Configure the minimum adjusted value for the node.

  [edit dynamic-profiles *profile-name* class-of-service traffic-control-profile *profile-name*]
  user@host# **set** adjust-minimum ***rate***

BEST PRACTICE: Ensure that the minimum adjusted value that you configure does not exceed the shaping rate and is not lower than the configured transmit rate.

Related Documentation

- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 117
- Configuring the Minimum Adjusted Shaping Rate on Scheduler Nodes for Subscribers on page 37
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3

Enabling Shaping-Rate Adjustments for Subscriber Local Loops

You can enhance a CoS implementation by enabling an MX Series 3D Universal Edge Router to adjust the hierarchical CoS policy shaping rate configured for static interface sets that consist of two or more VLANs and represent subscriber local loops. Whenever the digital subscriber line access multiplexer (DSLAM) resynchronizes its data transmission rate to a digital subscriber line (DSL), the router adjusts the shaping rate for the associated subscriber interface so that the maximum bandwidth allocation cannot exceed the current data rate for the associated subscriber local loop. This feature ensures that data transmission rate adjustments by the DSLAM do not cause bandwidth contention at the subscriber's residential gateway.

This topic includes the following tasks:

- Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops on page 40
- Configuring the Logical Interfaces That Compose the Static Logical Interface Sets on page 41
- Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops on page 42
- Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops on page 44

Configuring Static Logical Interface Sets to Serve as CoS Hierarchical Scheduler Nodes for Subscriber Loops

To configure a logical interface set, begin by including the **interface-set** statement with the *interface-set-name* option at the **[edit interfaces]** hierarchy level.

An interface set is composed of two or more logical interfaces on the same physical interface. Each logical interface in an interface set corresponds to an individual subscriber

service, such as voice, video, or data. To specify either a list of logical unit numbers or the single outer VLAN tag used to identify the logical interfaces that compose the interface set, include statements at the **[edit interfaces interface-set *interface-set-name*]** hierarchy level:

- For an interface set composed of a list of logical interfaces identified by an inner VLAN tag on Ethernet frames (called the customer VLAN, or C-VLAN, tag), you must specify each logical interface by including the **unit** statement with the ***logical-unit-number*** option.

```
[edit]
interfaces {
    interface-set interface-set-name {
        interface ethernet-interface-name { # EQ DPC port
            unit logical-unit-number;
            unit logical-unit-number;
            ...
        }
        ...
    }
}
```

- For an interface set composed of a set of VLANs grouped at the DSLAM and identified by the same service VLAN (S-VLAN) tag), you must specify the S-VLAN tag as the outer VLAN tag for each VLAN by including the **vlan-tags-outer** statement with the ***vlan-tag*** option.

```
[edit]
interfaces {
    interface-set interface-set-name {
        interface ethernet-interface-name { # EQ DPC port
            vlan-tags-outer vlan-tag; # Identify the DSLAM
        }
        ...
    }
}
```

For more information, see *Configuring Hierarchical Schedulers for CoS*.

Configuring the Logical Interfaces That Compose the Static Logical Interface Sets

Each underlying physical interface must be configured to operate in hierarchical scheduler mode and to support stacked VLAN tagging on all logical interfaces. To configure, include the **hierarchical-scheduler** statement and the **stacked-vlan-tagging** statement at the **[edit interfaces *ethernet-interface-name*]** hierarchy level.

To associate the individual logical interfaces of an interface set with specific subscriber services provided by the subscriber local loop, bind an S-VLAN tag and a C-VLAN tag to each logical interface that belongs to a scheduler node that represents a subscriber local loop. Ethernet frames sent from the logical interfaces contain an outer VLAN tag that identifies a DSLAM and an inner VLAN tag that identifies a subscriber port on the DSLAM. To configure, include the **vlan-tags** statement at each logical interface:

```
[edit]
interfaces {
```

```
    ethernet-interface-name { # EQ DPC port underlying an interface set
      hierarchical-scheduler;
      stacked-vlan-tagging; # Support 802.1Q VLAN dual-tagged frames
      unit logical-unit-number { # Bind S-VLAN and C-VLAN tags to logical interface
        vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
      }
      ...
    }
  }
```

For more information, see *802.1Q VLANs Overview*.

Configuring Hierarchical CoS on the Static Logical Interface Sets That Serve as Hierarchical Scheduler Nodes for Subscriber Local Loops

To configure hierarchical CoS on the static logical interface set that serves as the hierarchical scheduler node for a subscriber local loop:

1. For each scheduler node that represents a subscriber local loop, configure an initial shaping rate.

   NOTE: The CoS shaping-rate feature is supported only for scheduler nodes with a configured shaping rate. The initial shaping rate must be configured by applying a traffic-control profile that includes the **shaping-rate** statement. Specify the initial shaping rate as a peak rate, in bits per second (bps), and not as a percentage. Other methods of configuring a shaping rate are not supported with this feature.

   - To enable traffic heading downstream (from the router to the DSLAM) to be gathered into an interface set, include the **interface-set** statement and define the logical interface set name as the ***interface-set-name*** option at the **[edit class-of-service interfaces]** hierarchy level.
   - To apply output traffic scheduling and shaping parameters at the logical interface set level (rather than at the logical unit level), include the **output-traffic-control-profile** statement and specify the name of a traffic-control profile as the ***profile-name*** option at the **[edit class-of-service interfaces interface-set *interface-set-name*]** hierarchy level.

   To configure, include the following statements:

```
interfaces { # Configure interface-specific CoS for incoming packets
  interface-set interface-set-name { # Configure a hierarchical scheduler
    output-traffic-control-profile tc-profile-name; # Level 3 scheduler node
  }
  ...
}
traffic-control-profiles { # Define traffic-control profiles
  tc-profile-name { # Specify a scheduler map and traffic-shaping parameters
    scheduler-map map-name;
    shaping-rate rate; # This is the “configured shaping rate”
    guaranteed-rate (percent percentage | rate);
    delay-buffer-rate (percent percentage | rate);
```

```
    }
    ...
}
```

You can include the statements at the following hierarchy levels:

- **[edit** class-of-service**]**
- **[edit** *dynamic-profiles profile-name* class-of-service**]**

2. Configure the scheduler maps referenced in the traffic-control profiles applied to the interface sets, the schedulers referenced in those scheduler maps, and the drop profiles referenced in those schedulers.

   - A scheduler map establishes the traffic output queues (forwarding classes) for a scheduler node and associates each queue with a specific scheduler map.
   - A scheduler defines queue properties (transmit rate, buffer size, priority, and drop profile) that specify how traffic is treated in the output queue.
   - A drop profile specifies how aggressively the MX Series router drops packets that are managed by a particular scheduler by defining either a segmented or interpolated graph that maps output queue fullness to packet drop probability.

   To configure, include the statements at the static **[edit** class-of-service**]** hierarchy level:

```
[edit]
class-of-service {
  scheduler-maps { # Assign queuing characteristics to output queues
    map-name { # Map output queues to
      forwarding-class class-name scheduler scheduler-name;
      forwarding-class class-name scheduler scheduler-name;
      ...
    }
    ...
  }
  schedulers { # Define queuing characteristics
    scheduler-name { # Specify queuing and buffer management
      transmit-rate transmit-rate-option;
      buffer-size buffer-size-option;
      priority priority-level;
      drop-profile-map loss-priority loss-priority-option protocol any drop-profile
        drop-profile-name;
      ...
    }
  }
  drop-profiles { # Define random early detection (RED) for the delay buffer
    drop-profile–name { # Specify how to drop packets from an output queue
      drop-profile-name { # Map a queue fullness to a drop probability
        fill-level percentage drop-probability percentage; # Option 1: segmented
        fill-level percentage drop-probability percentage;
        ...
      }
        interpolate { # Option 2: interpolated
          drop-probability [ values ];
          fill-level [ values ];
```

```
            }
          }
        ...
      }
    }
```

For more information about configuring scheduler maps, schedulers, and drop profiles, see *CoS Inputs and Outputs Overview*.

Configuring ANCP Functionality That Supports and Drives Shaping-Rate Adjustments for Subscriber Local Loops

To configure the Access Node Control Protocol (ANCP) functionality that supports and drives the shaping-rate adjustments for subscriber local loops:

- Enable ANCP to monitor subscriber local loop rates at the DSLAMs and communicate this information to CoS.
- Configure each DSLAM as an ANCP neighbor of the router so that TCP connections can be established between the router and each DSLAM.
- Identify the subscriber interface sets whose traffic is monitored and shaped by ANCP, and associate those interface sets with the corresponding identifiers configured on the access node (DSLAM) to uniquely identify the subscriber local loops within the access network.

  ANCP uses this information to build a mapping of subscribers to subscriber interfaces. When ANCP receives port management messages from a DSLAM or other access node, it uses the access identifier contained in the message to determine which hierarchical scheduler node corresponds to the subscriber.

To configure, include statements at the **[edit *protocols ancp*]** hierarchy level:

```
[edit]
protocols {
  ancp {
    qos-adjust; # Enable ANCP to monitor and adjust CoS shaping rates
    neighbor ip-address; # Configure each DSLAM as an ANCP neighbor
    ...
    interfaces { # Identify subscribers for which ANCP can adjust shaping rates
      interface-set {
        interface-set-name {
          access-identifier identifier-string; # DSLAM ID for the local loop
        }
      }
      ...
    }
    ...
  }
  ...
}
```

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 33

- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 5
- Verifying the Configuration of ANCP for Shaping-Rate Adjustments on page 119
- Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops on page 118
- Disabling Shaping-Rate Adjustments for Subscriber Local Loops on page 45
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 65

Disabling Shaping-Rate Adjustments for Subscriber Local Loops

To disable hierarchical CoS shaping-rate adjustments for subscriber local loops:

- Disable hierarchical CoS traffic-shaping adjustment by ANCP:

```
[edit protocols ancp]
user@host# delete qos-adjust
```

Traffic-shaping parameters for all subscriber local loops revert to their current configured values.

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 33
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 5
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 40
- Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops on page 65

Disabling Hierarchical Bandwidth Adjustment for Subscriber Interfaces with Reverse-OIF Mapping

You can disable hierarchical bandwidth adjustment for all subscriber interfaces with reverse OIF mapping enabled on a specified multicast interface. Reverse OIF mapping is used to determine the subscriber VLAN interface and the multicast traffic bandwidth on the interface.

To disable hierarchical bandwidth adjustment:

1. Specify that you want to access the subscriber interfaces with reverse-OIF mapping enabled.

   ```
   [edit routing-instances routing-instance routing-options multicast interface
     interface-name]
   user@host# edit reverse-oif-mapping
   ```

2. Disable hierarchical bandwidth adjustment for all subscriber interfaces on the interface.

   ```
   user@host# set no-qos-adjust
   ```

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3
- *Example: Configuring Multicast with Subscriber VLANs*

Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces

To configure the PPPoE vendor-specific tags feature in a dynamic profile:

NOTE: **When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured for shaping-rate and overhead-accounting statements at the [edit dynamic-profiles profile-name class-of-service traffic-control-profile] hierarchy level.**

1. (Optional) To configure the shaping rate based on access line information:

   [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]
   user@host# set vendor-specific-tags actual-data-rate-downstream

2. (Optional)To configure the overhead-accounting based on access-line information:

   [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]
   user@host# set vendor-specific-tags access-loop-encapsulation

Related Documentation

- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 17
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9

Reporting the Effective Shaping Rate for Subscribers

The Effective-Shaping-Rate VSA [26–177] provides the best estimate for a subscriber's downstream traffic rate for accounting purposes. The VSA is included in RADIUS Acct-Start, Acct-Stop, and Interim-Acct messages. The reported rate is the rate enforced on the L3, L2, or L1 node according to local policy. The value of the VSA varies depending on your configuration:

- Actual rate—When effective shaping rate reporting is enabled.
- Advisory rate—When the advisory rate is configured and effective shaping rate reporting is not enabled.
- Port speed—When the advisory rate is not configured and effective shaping rate reporting is not enabled.

When you disable reporting, the VSA reports either the advisory rate or port speed for both existing subscribers and new subscribers that log in after reporting is disabled.

To enable reporting of the actual downstream traffic rate:

- Enable reporting.

```
[edit chassis]
user@host1# set effective-shaping-rate
```

NOTE: When the traffic control profile for the subscriber specifies cell-mode, the effective shaping rate does not account for cell padding according to the encapsulation type. The rate includes the 48/53 cell tax.

Related Documentation

- Verifying the Effective Shaping Rate Reporting Configuration on page 120
- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- *Juniper Networks VSAs Supported by the AAA Service Framework*
- *AAA Accounting Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS*

Configuring Scheduler and Scheduler Map Sharing

The system generates unique identifiers (IDs) in dynamic profiles created for services. The generated unique IDs enable you to identify and configure separate parameter values with the same variable name. When applied to CoS, you can now configure scheduler and scheduler map sharing. In client-access profiles, schedulers and scheduler maps must use the unique ID format. If the client-access profile uses the unique ID format and you want to have either scheduler or scheduler map sharing for service activation, you must configure the service profile in unique ID format. Generating unique IDs based on schedulers and scheduler maps eliminates duplication and improves router performance and scalability. You can configure scheduler and scheduler map sharing by including the variables for CoS in the client access or service dynamic profile. All scheduler maps and schedulers must be in the unique ID format.

Before you configure variables for the client access or service dynamic profile:

- Create a basic dynamic profile.

  See *Configuring a Basic Dynamic Profile*.

To configure variables for the client access or service dynamic profile:

1. Configure the variables for the dynamic client access profile.

   ```
   [edit dynamic-profiles client-profile variables]
   user@host# set smap_data uid
   user@host# set data_sched uid
   ```

2. Configure the CoS parameters for the variables in the scheduler profile.

   ```
   [edit dynamic-profilesclient-profile class-of-service]
   user@host# edit schedulers “$data_sched”
   user@host# set transmit-rate percent 10
   user@host# set buffer-size remainder
   user@host# set priority low
   ```

3. Configure the CoS parameters for the variables in the scheduler maps profile.

   [edit dynamic-profiles *client-profile* class-of-service]
   user@host# **edit scheduler-maps “$smap_data”**
   user@host# **edit forwarding-class be scheduler “$data_sched”**

For example, you can configure scheduler maps and schedulers for a client access profile:

```
dynamic-profiles {
  cos-para {
    variables {
      data_smap uid;
      data_video_smap uid;
      voice_smap uid;
      data_sched uid;
      video_sched uid;
      voice_sched uid;
    }
    ...
    class-of-service {
      traffic-control-profiles {
        tcp1 {
          scheduler-map "$junos-cos-scheduler-map";
          shaping-rate "$junos-cos-shaping-rate";
          guaranteed-rate 10m;
          delay-buffer-rate "$junos-cos-delay-buffer-rate";
        }
      }
      interfaces {
        "$junos-interface-ifd-name" {
          unit "$junos-underlying-interface-unit" {
            output-traffic-control-profile tcp1;
          }
        }
      }
      scheduler-maps {
        "$data_smap" {
          forwarding-class be scheduler "$data_sched";
        }
        "$data_video_smap" {
          forwarding-class be scheduler "$data_sched";
          forwarding-class af scheduler "$video_sched";
        }
        “$voice_smap” {
          forwarding-class ef scheduler “$voice_sched”;
        }
      }
      schedulers {
        "$data_sched" {
          transmit-rate "$junos-cos-scheduler-tx";
          inactive: buffer-size percent "$junos-cos-scheduler-bs";
          priority "$junos-cos-scheduler-pri";
        }
        "$video_sched" {
          transmit-rate "$junos-cos-scheduler-tx";
          inactive: buffer-size percent "$junos-cos-scheduler-bs";
```

```
            priority "$junos-cos-scheduler-pri";
        }
        “$voice_sched” {
            transmit-rate percent 10;
            buffer-size remainder;;
            priority low;
        }
    }
  }
 }
}
```

Related Documentation

- *Access Profiles and Service Profiles Overview*
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27

CHAPTER 7

Configuration Tasks for Traffic Distribution

Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

You can configure the overhead accounting feature to shape downstream traffic based on either frames or cells.

You can also account for the different byte sizes per encapsulation by configuring a byte adjustment value for the shaping mode.

This feature is supported on MPCs on MX Series routers.

To configure the overhead accounting feature in a dynamic profile:

1. Do one of the following to configure the shaping mode:
   - Specify the shaping mode.

     Frame shaping mode is enabled by default.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*
     user@host#set overhead-accounting **(frame-mode | cell-mode)**
   - Configure a variable for the shaping mode.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*
     user@host#set overhead-accounting **$junos-cos-shaping-mode**
2. (Optional) Do one of the following to configure the byte adjustment value:
   - Specify a byte adjustment value.

     [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*

```
user@host#set overhead-accounting bytes byte-value
```

- Configure a variable for the byte adjustment.

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles
  profile-name
user@host#set overhead-accounting bytes $junos-cos-byte-adjust
```

BEST PRACTICE: We recommend that you specify a byte adjustment value that represents the difference between the customer premise equipment (CPE) protocol overhead and B-RAS protocol overhead.

The available range is –120 through 124 bytes. The system rounds up the byte adjustment value to the nearest multiple of 4. For example, a value of 6 is rounded to 8, and a value of -10 is rounded to -8.

Related Documentation

- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 61
- Verifying the Scheduling and Shaping Configuration for Subscriber Access on page 117

Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces

Service providers often used tiered services that must utilize excess bandwidth as traffic patterns vary. By default, excess bandwidth between a configured guaranteed rate and shaping rate is shared equally among all queues with the same excess priority value, which might not be optimal for all subscribers to a service.

This feature is supported for MIC and MPC interfaces on MX Series routers.

To configure parameters to manage excess bandwidth for subscriber interfaces:

1. Configure the parameters for the interface.
   a. Configure the guaranteed and shaping rates.
      i. Configure the guaranteed rate:

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles
  profile-name]
user@host# set guaranteed-rate(rate | $junos-cos-guaranteed-rate) <burst-size
  (bytes | $junos-cos-guaranteed-rate-burst)>
```

      ii. Configure the shaping rate:

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles
  profile-name]
user@host# set shaping-rate (rate | $junos-cos-shaping-rate) <burst-size (bytes
  | $junos-cos-shaping-rate-burst)>
```

TIP: On MPC/MIC interfaces, the guaranteed rate and the shaping rate share the value specified for the burst size. If the guaranteed rate has a burst size specified, it is used for the shaping rate; if the shaping rate has a burst size specified, it is used for the guaranteed rate. If you have specified a burst for both rates, the system uses the lesser of the two values.

b. Configure a rate for excess bandwidth.

You can configure an excess rate for all priorities of traffic:

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles
  profile-name]
user@host# set excess-rate (percent percentage | $junos-cos-excess-rate) |
  proportion value )
```

Optionally, you can configure an excess rate specifically for high- and low-priority traffic. When you configure the **excess-rate** statement for an interface, you cannot also configure the **excess-rate-low** and **excess-rate-high** statements.

```
[edit dynamic-profiles profile-name class-of-service traffic-control-profiles
  profile-name]
user@host# set excess-rate-high(percent percentage |
  $junos-cos-excess-rate-high) | proportion value )
user@host# set excess-rate-low (percent percentage | $junos-cos-excess-rate-low)
  | proportion value )
```

BEST PRACTICE: We recommend that you configure either a percentage or a proportion of the excess bandwidth for all schedulers with the same parent in the hierarchy. For example, if you configure interface 1.1 with twenty percent of the excess bandwidth, configure interface 1.2 with eighty percent of the excess bandwidth.

2. (Optional) Configure parameters for the queue.

a. Configure the shaping rate.

```
[edit dynamic-profiles profile-name class-of-service scheduler scheduler-name]
user@host#set shaping-rate (rate | $junos-cos-scheduler-shaping-rate) <burst-size
  bytes>
```

b. Configure the excess rate.

```
[edit dynamic-profiles profile-name class-of-service scheduler scheduler-name]
user@host#set excess-rate (percent percentage | percent
  $junos-cos-scheduler-excess-rate)
```

c. (Optional) Configure the priority of excess bandwidth for the queue.

```
[edit dynamic-profiles profile-name class-of-service scheduler scheduler-name]
user@host#set excess-priority (low | high | $junos-cos-scheduler-excess-priority
  | none)
```

TIP:

For queues, you cannot configure the excess rate or excess priority in these cases:

- When the **transmit-rate exact** statement is configured. In this case, the shaping rate is equal to the transmit rate and the queue does not operate in the excess region.
- When the scheduling priority is configured as **strict-high**. In this case, the queue gets all available bandwidth and never operates in the excess region.

By default, when traffic exceeds the shaping or guaranteed rates, the system demotes traffic configured with high or medium priority. To disable priority demotion, specify the **none** option. You cannot configure this option for queues configured with **transmit-rate** expressed as a percent and when the parent's guaranteed rate is set to zero.

Related Documentation

- For hardware requirements and configuration guidelines, see Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27

Providing Accurate Scheduling for a Demux Subscriber Interface of Aggregated Ethernet Links

Unlike VLAN subscriber interfaces, enabling link protection is not required for configuring hierarchical CoS on demux interfaces. Instead, we recommend that you enable targeted distribution on the demux interface to provide accurate scheduling for the aggregated Ethernet links.

Before you begin, configure the subscriber interface with aggregated Ethernet:

- For static and dynamic IP demux interfaces, see *Configuring a Static or Dynamic IP Demux Subscriber Interface over Aggregated Ethernet*.
- For static and dynamic VLAN demux interfaces, see *Configuring a Static or Dynamic VLAN Demux Subscriber Interface over Aggregated Ethernet*.

To provide accurate scheduling for a demux subscriber interface of aggregated Ethernet links:

1. Enable targeted distribution for the demux interface.

   See *Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces*.

2. Enable hierarchical scheduling on the link aggregation bundle.

   See *Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links*.

3. (Optional) Enable module redundancy to ensure that CoS resources are provisioned for the aggregated Ethernet links if a module or a link fails. By default, link redundancy is supported.

   See *Configuring Link and Module Redundancy for Demux Subscribers in an Aggregated Ethernet Interface*.

4. (Optional) Configure rebalancing periodically or manually for the subscribers. See *Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface*.

5. Attach static or dynamic traffic shaping and scheduling parameters at the aggregated Ethernet logical interface or its underlying physical interface. See:

   - *Configuring Traffic Scheduling and Shaping for Subscriber Access*
   - *Configuring Schedulers in a Dynamic Profile for Subscriber Access*
   - *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*
   - *Applying Minimal Shaping and Scheduling to Remaining Subscriber Traffic*

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface on page 118

CHAPTER 8

Configuration Tasks for Scaling

Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces

This topic describes how to manage dedicated and remaining queues for static and dynamic subscriber interfaces configured in dynamic profiles.

You manage queues at the chassis and physical port level in the static configuration hierarchies, then configure dynamic scheduling and shaping parameters for the subscriber interfaces in the dynamic profile.

Configuring the Maximum Number of Queues for MIC and MPC Interfaces

30-Gigabit Ethernet Queuing MPCs and 60-Gigabit Ethernet Queuing and Enhanced Queuing MPCs support a dedicated number of queues when configured for hierarchical scheduling and per-unit scheduling configurations.

To scale the number of subscriber interfaces per queue, you can modify the number of queues supported on the MIC.

To configure the number of queues:

1. Specify that you want to configure the MIC.

   user@host# **edit chassis fpc** *slot-number* **pic** *pic-number*

2. Configure the number of queues.

   [edit chassis fpc *slot-number* pic *pic-number*]
   user@host# **set**max-queues-per-interface **(8 | 4)**

Configuring Remaining Common Queues on MIC and MPC Interfaces

30-Gigabit Ethernet Queuing MPCs and 60-Gigabit Ethernet Queuing and Enhanced Queuing MPCs support a dedicated set of queues when configured with hierarchical scheduling.

When the number of dedicated queues is reached on the module, there can be queues remaining. Traffic from these logical interfaces are considered unclassified and attached to a common set of queues that are shared by all subsequent logical interfaces.

You can configure traffic shaping and scheduling resources for the remaining queues by attaching a special traffic-control profile to the interface. This feature enables you to provide the same shaping and scheduling to remaining queues as the dedicated queues.

To configure the remaining queues on a MIC or MPC interface:

1. Configure CoS parameters in a traffic-control profile.

   [edit class-of-service]
   user@host# **edit traffic-control-profiles *profile-name***

2. Enable hierarchical scheduling for the interface.

   [edit interfaces *interface-name*]
   user@host# **set hierarchical-scheduler**

3. Attach the traffic control profiles for the dedicated and remaining queues to the port on which you enabled hierarchical scheduling.

   To provide the same shaping and scheduling parameters to dedicated and remaining queues, reference the same traffic-control profile.

   a. Attach the traffic-control profile for the dedicated queues on the interface.

      [edit class-of-service interfaces *interface-name*]
      user@host# **set output-traffic-control-profile *profile-name***

   b. Attach the traffic-control profile for the remaining queues on the interface.

      [edit class-of-service interfaces *interface-name*]
      user@host# **set output-traffic-control-profile-remaining *profile-name***

Related Documentation

- Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces on page 117
- Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview on page 19
- *Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*

CHAPTER 9

Examples

Example: Configuring Initial CoS Parameters Dynamically Obtained from RADIUS

The following configuration is an example of a client dynamic profile in which initial CoS parameters are dynamically obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is applied.

For this example, assume that the RADIUS authentication server has been configured with traffic-shaping parameters (at Juniper Networks VSA 26-108) and CoS scheduling and queuing parameters (at Juniper Networks VSA 26–146).

The subscriber interface is a single-unit static gigabit Ethernet VLAN interface on an EQ DPC port:

```
[edit]
interfaces {
   ge-9/0/3 {
      hierarchical-scheduler;
      vlan-tagging;
      unit 100 {
         vlan-id 100;
         family inet {
            address 192.168.32.2/24;
         }
      }
   }
}
```

The client dynamic profile **residential_silver** attaches the traffic-control profile **tcp_1** to the subscriber interface that is defined in the dynamic profile using the **$junos-interface-ifd-name** predefined variable.

```
[edit]
dynamic-profiles {
```

```
    residential_silver {
      interfaces {
        “$junos-interface-ifd-name” {
          unit “$junos-underlying-interface-unit” {
            family inet;
          }
        }
      }
      class-of-service {
        interfaces {
          “$junos-interface-ifd-name” {
            unit “$junos-underlying-interface-unit” {
              output-traffic-control-profile tcp_1;
            }
          }
        }
      }
    }
  }
```

The traffic-control profile **tcp_1**, references Junos OS predefined variables to obtain a scheduler-map name and traffic-shaping parameter values from RADIUS when a subscriber logs in. For this example, assume that the RADIUS server replaces the Junos OS predefined variable **$junos-cos-scheduler-map** scheduler-map name **business_smap_1**. The scheduler map **business_smap_1** is configured in the client dynamic profile:

```
[edit]
dynamic-profiles {
  residential_silver {
    class-of-service {
      traffic-control-profiles {
        tcp_1 {
          scheduler-map “$junos-cos-scheduler-map”; # ’business_smap_1’
          shaping-rate "$junos-cos-shaping-rate";
          guaranteed-rate "$junos-cos-guaranteed-rate";
          delay-buffer-rate "$junos-cos-delay-buffer-rate";
        }
      }
      scheduler-maps {
        business_smap_1 {
          forwarding-class best-effort scheduler be_sched;
          forwarding-class ef scheduler home_sched
        }
      }
    }
  }
}
```

A scheduler definition references Junos OS predefined variables to obtain scheduler configurations from RADIUS when a subscriber logs in. For this example, assume that the RADIUS server provides scheduler configurations for schedulers named **be_sched** and **home_sched**, which are included in the scheduler map **business_smap_1**:

```
[edit]
dynamic-profiles {
  residential_silver {
    class-of-service {
```



```
        schedulers {
          “$junos-cos-scheduler” { # ’be_sched’ and ’home_sched’
            transmit-rate "$junos-cos-scheduler-tx";
            buffer-size "$junos-cos-scheduler-bs";
            priority "$junos-cos-scheduler-pri";
            drop-profile-map loss-priority low protocol any drop-profile
              “$junos-cos-scheduler-dropfile-low“;
            drop-profile-map loss-priority medium-low protocol any drop-profile
              “$junos-cos-scheduler-dropfile-medium-low“;
            drop-profile-map loss-priority medium-high protocol any drop-profile
              “$junos-cos-scheduler-dropfile-medium-high“;
            drop-profile-map loss-priority high protocol any drop-profile
              “$junos-cos-scheduler-dropfile-high“;
          }
        }
      }
    }
  }
```

Static configurations for CoS consist of configurations for the forwarding classes used in the scheduler map **business_smap_1** and configurations for drop-profile names provided by RADIUS for as part of the scheduler configurations provided (for **be_sched** and **home_sched**) when a subscriber logs in:

```
[edit]
  class-of-service {
    forwarding-classes {
      queue 0 best-effort;
      queue 1 ef;
    }
    drop-profiles {
      . . . configurations_for_drop_profile_names_provided_by_RADIUS . . .
    }
  }
}
```

Related Documentation

- *Subscriber Activation and Service Management in an Access Network*
- *Dynamic Profiles Overview*
- *Dynamic Variables Overview*
- *Junos OS Predefined Variables*
- *Subscriber Interfaces That Provide Initial CoS Parameters Dynamically Obtained from RADIUS*
- *Configuring Initial CoS Parameters Dynamically Obtained from RADIUS*

Example: Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates

This topic describes two scenarios for which you can configure dynamic shaping parameters to account for packet overhead in a downstream network.

The RADIUS administrator supplies the initial values on the RADIUS server, and the service activation is performed at subscriber login.

Figure 8 on page 62 shows the sample network that the examples reference.

Figure 8: Sample Network Topology for Downstream Traffic

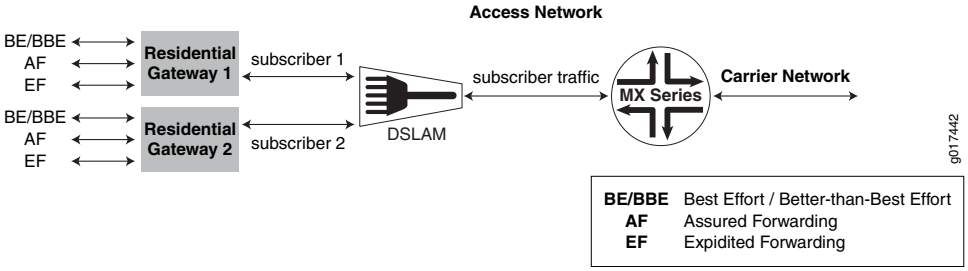

Managing Traffic with Different Encapsulations

In this example, the MX Series router shown in Figure 8 on page 62 sends stacked VLAN frames to the DSLAM, and the DSLAM sends single-tagged VLAN frames to the residential gateway.

To accurately shape traffic at the residential gateway, the MX Series router must account for the different frame sizes. The difference between the stacked VLAN (S-VLAN) frames sent by the router and the single-tagged VLAN frames received at the residential gateway is a 4-byte VLAN tag. The residential gateway receives frames that are 4 bytes less.

To account for the different frame sizes, you configure the frame shaping mode with -4 byte adjustment:

1. Configure the traffic shaping parameters in the dynamic profile and attach them to the interface.

   Enabling the overhead accounting feature affects the resulting shaping rate, guaranteed rate, and excess rate parameters, if they are configured.

```
[edit]
dynamic-profiles {
  ethernet-downstream-network {
  interfaces {
    $junos-interface-ifd-name {
      unit $junos-underlying-interface-unit {
        family inet;
      }
    }
  }
  class-of-service {
    traffic-control-profiles {
      tcp-example-overhead-accounting-frame-mode {
        excess-rate percent $junos-cos-excess-rate
        guaranteed-rate $junos-cos-guaranteed-rate
        overhead-accounting $junos-cos-shaping-mode bytes $junos-cos-byte-adjust
        shaping-rate $junos-cos-shaping-rate;
      }
    }
```

```
        interfaces {
          $junos-interface-ifd-name {
            unit "$junos-underlying-interface-unit" {
            output-traffic-control-profile tcp1;
            }
          }
        }
      }
      }
    }
```

Table 11 on page 63 lists the initial values defined by the RADIUS administrator for the shaping rates.

Table 11: Initial Shaping Values at Subscriber Login

Predefined Variable	RADIUS Tag	Value
$junos-cos-shaping-rate	T02	10m
$junos-cos-guaranteed-rate	T03	2m
$junos-cos-excess-rate	T05	50
$junos-cos-shaping-mode	T07	frame-mode
$junos-cos-byte-adjust	T08	–4

2. Verify the adjusted rates.

```
user@host#show class-of-service traffic-control-profile
Traffic control profile: tcp-example-overhead-accounting-frame-mode, Index:
61785
Excess rate 50
Shaping rate: 10000000
Guaranteed rate: 2000000
Overhead accounting mode: Frame Mode
Overhead bytes: –4
```

Managing Downstream Cell-Based Traffic

In this example, the DSLAM and residential gateway shown in Figure 8 on page 62 are connected through an ATM cell-based network. The MX Series router sends Ethernet frames to the DSLAM, and the DSLAM sends ATM cells to the residential gateway.

To accurately shape traffic at the residential gateway, the MX Series router must account for the different physical network characteristics.

The administrator does not need to configure a byte adjustment value to account for the downstream ATM network, but has the option of configuring a byte adjustment value to account for different encapsulations or decapsulations.

To account for the different frame sizes, configure cell shaping mode:

1. Configure the traffic shaping parameters in the dynamic profile and attach them to the interface.

   Enabling the overhead accounting feature affects the resulting shaping rate, guaranteed rate, and excess rate parameters, if they are configured.

```
[edit]
dynamic-profiles {
  atm-downstream-network {
  interfaces {
    $junos-interface-ifd-name {
      unit $junos-underlying-interface-unit {
        family inet;
      }
    }
  }
  class-of-service {
    traffic-control-profiles {
      tcp-example-overhead-accounting-cell-mode {
        excess-rate percent $junos-cos-excess-rate
        guaranteed-rate $junos-cos-guaranteed-rate
        overhead-accounting $junos-cos-shaping-mode
        shaping-rate $junos-cos-shaping-rate
      }
    }
    interfaces {
      $junos-interface-ifd-name {
        unit "$junos-underlying-interface-unit" {
        output-traffic-control-profile tcp1;
        }
      }
    }
  }
  }
}
```

   Table 12 on page 64 lists the initial values defined by the RADIUS administrator for the shaping rates.

Table 12: Initial Shaping Values at Subscriber Login

Predefined Variable	RADIUS Tag	Value
$junos-cos-shaping-rate	T02	10m
$junos-cos-guaranteed-rate	T03	2m
$junos-cos-excess-rate	T05	50
$junos-cos-shaping-mode	T07	cell-mode

2. Verify the adjusted rates.

```
user@host#show class-of-service traffic-control-profile
```

```
Traffic control profile: tcp-example-overhead-accounting-cell-mode, Index:
61785
Shaping rate: 10000000
Excess rate 50
Guaranteed rate: 2000000
Overhead accounting Cell Mode
Overhead bytes: 0
```

To account for ATM segmentation, the MX Series router adjusts all of the rates by 48/53 to account for ATM AAL5 encapsulation. In addition, the router accounts for cell padding, and internally adjusts each frame by 8 bytes to account for the ATM trailer.

Related Documentation

- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 51

Example: Configuring Hierarchical CoS Shaping-Rate Adjustments for Subscriber Local Loops

This example shows how you can enable shaping-rate adjustments for static logical interface sets that represent subscriber local loops:

1. Configure static logical interface sets to serve as CoS hierarchical scheduler nodes for subscriber local loops.

   This example uses a single scheduler node that represents two subscriber local loops. The scheduler node is a static logical interface composed of two logical interfaces. The underlying physical interface is port **0** on a Gigabit Ethernet EQ DPC in slot 4, PIC 0:

```
[edit]
interfaces {
  interface-set ifset-of-logical-interfaces {
    interface ge-4/0/0 {
      unit 1;
      unit 2;
    }
  }
  ge-4/0/0 {
    description “access interface ge-4/0/0”;
    hierarchical-scheduler;
    stacked-vlan-tagging;
    unit 1 {
      description “DSL type ADSL1 = 0x01”;
      proxy-arp;
      vlan-tags outer 1 inner 1; # S-VLAN tag is ’1’ and C-VLAN tag is ’1’
      family inet { # Specify a secondary loopback address
        unnumbered-address lo0.0 preferred-source-address 192.168.7.3;
      }
    }
    unit 2 {
      description “DSL type ADSL1 = 0x01”;
      proxy-arp;
      vlan-tags outer 1 inner 2; # S-VLAN tag is ’1’ and C-VLAN tag is ’2’
```

```
        family inet { # Specify a secondary loopback address
          unnumbered-address lo0.0 preferred-source-address 192.168.7.4;
        }
      }
    }
  }
```

2. Begin configuring hierarchical CoS on the static logical interface set that serves as the hierarchical scheduler node for the group of subscriber local loops.

```
[edit]
class-of-service {
  interfaces {
    interface-set ifset-of-logical-interfaces {
      output-traffic-control-profile tcp-premium-with-4–queues;
    }
  }
}
```

3. Configure the traffic-control profiles that can be applied to the scheduler node:

```
[edit]
class-of-service {
  traffic-control-profiles {
    tcp-basic-rate { # Specify a scheduler map and traffic controls
      shaping-rate 10m;
    }
    tcp-premium-with-4-queues { # Specify a scheduler map and traffic controls
      scheduler-map smap-premium-4q;
      shaping-rate 20m;
      guaranteed-rate 10m;
      delay-buffer-rate 5m;
    }
  }
}
```

In this example, the **tcp-premium-with-4-queues** traffic-control profile is applied to the interface set. The other profile provides a lower shaping rate and no guaranteed rate.

4. Configure the scheduler map **smap-premium-4q** that is referenced in the traffic-control profile for the scheduler node:

```
[edit]
class-of-service {
  scheduler-maps { # Define the queues that comprise each scheduler node
    smap-premium-4q { # Map each queue in the scheduler node to a scheduler
      forwarding-class be scheduler be_sch;
      forwarding-class af scheduler af_sch;
      forwarding-class ef scheduler ef_sch;
      forwarding-class nc scheduler nc_sch;
    }
  }
}
```

5. Configure the four schedulers (referenced in the scheduler map) that define the four output queues for the scheduler node:

```
[edit]
class-of-service {
  schedulers { # Define scheduling characteristics of each queue
    be_sch { # Transmit rate and buffer management parameters
      transmit-rate percent 10;
      buffer-size remainder;
      priority low;
    }
    ef_sch { # Transmit rate and buffer management parameters
      ...
    }
    af_sch { # Transmit rate and buffer management parameters
      ...
    }
    nc_sch { # Transmit rate and buffer management parameters
      ...
    }
  }
}
```

6. Enable ANCP to communicate with the DSLAM to adjust the CoS shaping rate for the scheduler node.

   You must enable the ANCP feature for performing CoS traffic shaping adjustments, configure the DSLAM as an ANCP neighbor, and specify the DSLAM-assigned identifier for the subscriber local loop represented by the scheduler node:

```
[edit]
protocols {
  ancp {
    qos-adjust; # Enable ANCP to adjust CoS shaping rates
    neighbor 10.2.3.4; # Configure the DSLAM as an ANCP neighbor
    interfaces { # Identify subscribers for which ANCP can adjust shaping rates
      interface-set {
        ifset-of-logical-interfaces {
          access-identifier “dslam port 2/3”; # DSLAM ID for the local loop
        }
      }
    }
  }
}
```

NOTE: If ANCP is not yet enabled, the process starts when you commit a configuration that contains the **protocols ancp** stanza.

7. You can display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical CoS, issue the **show class-of-service interface-set** operational command.

Related Documentation

- Hierarchical CoS Shaping-Rate Adjustments Overview on page 3
- Shaping Rate Adjustments for Subscriber Local Loops Overview on page 5
- Guidelines for Configuring Shaping-Rate Adjustments for Subscriber Local Loops on page 33
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 40

CHAPTER 10

Configuration Statements

[edit dynamic-profiles] Hierarchy Level

```
dynamic-profiles {
  profile-name {
    class-of-service {
      interfaces {
        interface-name {
          unit logical-unit-number {
            classifiers {
              type (classifier-name | default);
            }
            output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
            rewrite-rules {
              dscp (rewrite-name | default);
              dscp-ipv6 (rewrite-name | default);
              ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
              inet-precedence (rewrite-name | default);
              }
            }
          }
        }
      }
      scheduler-maps {
        map-name {
          forwarding-class class-name scheduler scheduler-name;
        }
      }
      schedulers {
        (scheduler-name) {
          buffer-size (percent percentage | remainder | temporal microseconds |
            $junos-cos-scheduler-bs);
          drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
            protocol (any | non-tcp | tcp) drop-profile (profile-name | predefined-variable);
          excess-priority (low | high | $junos-cos-scheduler-excess-priority);
          excess-rate (percent percentage  | percent $junos-cos-scheduler-excess-rate);
          overhead-accounting (shaping-mode) <bytes (byte-value>;
          priority (priority-level | $junos-cos-scheduler-priority);
          shaping-rate (rate | predefined-variable);
          transmit-rate (rate | percent percentage | remainder | percent percentage
            $junos-cos-scheduler-tx) <exact | rate-limit>;
        }
      }
      traffic-control-profiles profile-name {
```

```
        delay-buffer-rate (percent percentage | rate);
        excess-rate (percent percentage | proportion value | percent
          $junos-cos-excess-rate);
        guaranteed-rate (percent percentage | rate);
        overhead-accounting (shaping-mode) <bytes (byte-value>;
        scheduler-map map-name;
        shaping-rate (percent percentage | rate | predefined-variable);
      }
    }
    firewall {
      family family {
        fast-update-filter filter-name {
          interface-specific;
        match-order [match-order];
          term term-name {
            from {
              match-conditions;
            }
            then {
              action;
              action-modifiers;
            }
          only-at-create;
        filter filter-name {
        interface-specific;
          term term-name {
            from {
              match-conditions;
            }
            then {
              action;
              action-modifiers;
            }
          }
      policer policer-name {
        filter-specific;
        if-exceeding {
          (bandwidth-limit bps | bandwidth-percent percentage);
          burst-size-limit bytes;
        }
        logical-bandwidth-policer;
        logical-interface-policer;
        physical-interface-policer;
        then {
          policer-action;
        }
      }
      hierarchical-policer policer-name {
        aggregate {
          if-exceeding {
            bandwidth-limit-limit bps;
            burst-size-limit bytes;
          }
          then {
            policer-action;
          }
```

```
        }
        premium {
          if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
          }
          then {
            policer-action;
          }
        }
      }
      three-color-policer policer-name {
        action {
          loss-priority high then discard;
        }
        logical-interface-policer;
        single-rate {
          (color-aware | color-blind);
          committed-burst-size bytes;
          committed-information-rate bps;
          excess-burst-size bytes;
        }
        two-rate {
          (color-aware | color-blind);
          committed-burst-size bytes;
          committed-information-rate bps;
          peak-burst-size bytes;
          peak-information-rate bps;
          }
        }
      }
    }
    policy-options {
      prefix-listname {
        ip-addresses;
      }
    }
    interfaces {
      interface-name {
        unit logical-unit-number {
          family family {
            access-concentrator name;
            address address;
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
              adf {
                counter;
                input-precedence precedence;
                not-mandatory;
                output-precedence precedence;
                rule rule-value;
              }
              input filter-name {
                precedence precedence;
```

```
            shared-name filter-shared-name;
          }
          output filter-name {
            precedence precedence;
            shared-name filter-shared-name;
          }
        }
        max-sessions number;
        max-sessions-vsa-ignore;
        rpf-check {
          fail-filter filter-name;
          mode loose;
        }
        service {
          input {
            service-set service-set-name {
              service-filter filter-name;
            }
            post-service-filter filter-name;
          }
          output {
            service-set service-set-name {
              service-filter filter-name;
            }
          }
        }
        service-name-table table-name;
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
          maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
      }
      ppp-options {
        chap;
        pap;
      }
      vlan-id number;
    }
    vlan-tagging;
  }
  interface-set interface-set-name {
    interface interface-name {
      unit logical-unit-number;
    }
  }
  demux0 {
    unit logical-unit-number {
      demux-options {
        underlying-interface interface-name
      }
      demux-source {
        source-prefix;
      }
      family family {
        access-concentrator name;
        address address;
        direct-connect;
```

```
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
          input filter-name;
          output filter-name;
        }
        mac-validate (loose | strict):
        max-sessions number;
        max-sessions-vsa-ignore;
        service-name-table table-name;
        short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
          maximum-seconds>;
        unnumbered-address interface-name <preferred-source-address address>;
      }
    }
  }
  pp0 {
    unit logical-unit-number {
      keepalives interval seconds;
      no-keepalives;
      pppoe-options {
        underlying-interface interface-name;
        server;
      }
      ppp-options {
        authentication [ authentication-protocols ];
        chap {
          challenge-length minimum minimum-length maximum maximum-length;
        }
        pap;
      }
      family inet {
        unnumbered-address interface-name address;
        address address;
        service {
          input {
            service-set service-set-name {
              service-filter filter-name;
            }
            post-service-filter filter-name;
          }
          output {
            service-set service-set-name {
              service-filter filter-name;
            }
          }
        }
        filter {
          input filter-name {
            precedence precedence;
          }
          output filter-name {
            precedence precedence;
          }
        }
      }
```

```
    }
  }
}
protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-policy;
      immediate-leave
      no-accounting;
      promiscuous-mode;
      ssm-map ssm-map-name;
      static {
        group group {
          source source;
        }
      }
      version version;
    }
    mld {
      interface interface-name {
        disable;
        (accounting | no-accounting);
        group-policy;
        immediate-leave;
        oif-map;
        passive;
        ssm-map ssm-map-name;
        static {
          group multicast-group-address {
            exclude;
            group-count number;
            group-increment increment;
            source ip-address {
              source-count number;
              source-increment increment;
            }
          }
        }
        version version;
      }
    }
    router-advertisement {
      interface interface-name {
        current-hop-limit number;
        default-lifetime seconds;
        (managed-configuration | no-managed-configuration);
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
        (other-stateful-configuration | no-other-stateful-configuration);
        prefix prefix {
          (autonomous | no-autonomous);
          (on-link | no-on-link);
          preferred-lifetime seconds;
          valid-lifetime seconds;
```

```
            }
            reachable-time milliseconds;
            retransmit-timer milliseconds;
          }
        }
      }
    }
  }
  routing-instances routing-instance-name {
    interface interface-name;
    routing-options {
      access {
        route prefix {
          next-hop next-hop;
          metric route-cost;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
      multicast {
        interface interface-name {
          no-qos-adjust;
        }
      }
    }
    rib routing-table-name {
      access {
        route prefix {
          next-hop next-hop;
          metric route-cost;
          preference route-distance;
          tag route-tag;
        }
      }
      access-internal {
        route subscriber-ip-address {
          qualified-next-hop underlying-interface {
            mac-address address;
          }
        }
      }
    }
  }
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
```

```
            tag route-tag;
          }
        }
        access-internal {
          route subscriber-ip-address {
            qualified-next-hop underlying-interface {
              mac-address address;
            }
          }
        }
        multicast {
          interface interface-name {
            no-qos-adjust;
          }
        }
      }
      variables {
        variable-name {
          default-value default-value;
          equals expression;
          mandatory;
          uid;
          uid-reference;
        }
      }
    }
```

Related Documentation

- *Dynamic Profiles Overview*
- *CoS for Subscriber Access Overview*
- *Configuring a Basic Dynamic Profile*
- *Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Two-Color Policer Configuration Overview*
- *Three-Color Policer Configuration Overview*
- *Hierarchical Policer Configuration Overview*
- *Guidelines for Applying Traffic Policers*

[edit class-of-service] Hierarchy Level

This topic shows the complete configuration for class of service (CoS) statements for the **[edit class-of-service]** hierarchy level, listing all possible configuration statements and showing their level in the configuration hierarchy. When you are configuring the Junos OS, your current hierarchy level is shown in the banner on the line preceding the **user@host#** prompt.

```
[edit class-of-service]
adjustment-control-profiles {
  profile-name {
    application {
```

```
        ancp;
        radius-coa;
        pppoe-tags;
        }
      }
    }
  classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority level code-points [ aliases ] [ bit-patterns ];
      }
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) {
      alias-name bits;
    }
  }
  copy-plp-all;
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability [ values ];
        fill-level [ values ];
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-classes {
    class class-name queue-num queue-number priority (high | low);
    queue queue-number class-name priority (high | low) [ policing-priority (premium |
      normal) ];
  }
  forwarding-class-map  forwarding-class-map-name {
    class class-name queue-num queue-number [ restricted-queue queue-number ];
  }
  forwarding-policy {
    next-hop-map map-name {
      forwarding-class class-name {
        next-hop [ next-hop-name ];
        lsp-next-hop [ lsp-regular-expression ];
        non-lsp-next-hop;
        discard;
      }
    }
    class class-name {
      classification-override {
        forwarding-class class-name;
      }
    }
```

```
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  forwarding-class class-name;
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
    rewrite-rules;
  }
}
interfaces {
  interface-name {
    classifiers {
      dscp (classifier-name | default);
      ieee-802.1 (classifier-name | default)  vlan-tag (inner | outer | classifier-name);
      inet-precedence (classifier-name | default);
    }
    input-scheduler-map map-name;
    input-shaping-rate rate;
    irb {
      unit logical-unit-number {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          exp (classifier-name | default);
          ieee-802.1 (classifier-name | default)  vlan-tag (inner | outer | transparent);
          }
          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);
          }
        }
      }
    }
    output-forwarding-class-map forwarding-class-map-name;
    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 {
            (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) (classifier-name | default)
              family (mpls | inet);
          }
          forwarding-class class-name;
          fragmentation-map map-name;
          input-scheduler-map map-name;
          input-shaping-rate (percent percentage | rate);
          input-traffic-control-profile profile-name shared-instance instance-name;
          loss-priority-maps {
            frame-relay-de (name | default);
          }
          loss-priority-rewrites {
            frame-relay-de (name | default);
          }
          output-traffic-control-profile profile-name shared-instance instance-name;
          per-session-scheduler;
          rewrite-rules {
            dscp (rewrite-name | default)protocol protocol-types;
            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)protocol protocol-types;
          }
          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;
        }
      }
    }
    loss-priority-maps {
      frame-relay-de (Defining Loss Priority Maps)name {
      loss-priority levelcode-points [alias | bits ];
      }
    }
    loss-priority-rewrites {
      frame-relay-de (Defining Loss Priority Maps)name {
      loss-priority levelcode-point (alias | bits );
      }
    }
    restricted-queues {
      forwarding-class class-name queue queue-number;
    }
    rewrite-rules {
      (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) rewrite-name {
        import (rewrite-name | default);
        forwarding-class class-name {
          loss-priority level code-point (alias | bits);
```

```
        }
    }
}
routing-instances routing-instance-name {
    classifiers {
        exp (classifier-name | default);
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        buffer-size (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 | high);
        excess-rate percent percentage;
        excess-rate (percent percentage | proportion value);
        priority priority-level;
        transmit-rate (rate | percent percentage | remainder) <exact | rate-limit>;
    }
}
system-defaults {
    classifiers (classifier-name | exp)
traffic-control-profiles profile-name {
    delay-buffer-rate (percent percentage | rate);
    excess-rate (percent percentage | proportion value);
    guaranteed-rate (percent percentage | rate);
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    scheduler-map map-name;
    shaping-rate (percent percentage | 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 {
        to-code-point value from-code-points (* | [ values ]);
    }
}
tri-color;
```

On Juniper Networks MX Series 3D Universal Edge Routers with Enhanced Queuing DPCs, you can configure the following CoS statements at the **[edit class-of-service interfaces]** hierarchy level:

```
interface-set interface-set-name {
    excess-bandwidth-share (proportional value | equal);
    internal-node;
    traffic-control-profiles profile-name;
    output-traffic-control-profile-remaining profile-name;
}
```

adjustment-control-profiles

Syntax

```
adjustment-control-profiles {
  profile-name {
    application {
      ancp;
      radius-coa;
      pppoe-tags;
      }
    }
  }
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 13.2.

Description Configure the CoS adjustment control profile.

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

The remaining statements are explained separately.

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- application (Adjustment Control Profiles) on page 86
- overhead-accounting (Dynamic Traffic Shaping) on page 103

adjust-minimum (Dynamic Shaping and Scheduling)

Syntax adjust-minimum (*rate* | $junos-cos-adjust-minimum);

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

Release Information Statement introduced in Junos OS Release 11.4.

Description For adjustments performed by the ANCP or multicast applications on EQ DPCs and MPC/MIC interfaces, specify the minimum shaping rate for an adjusted scheduler node. The node is associated with a traffic-control profile.

For adjustments performed by the multicast application on MPC/MIC interfaces, specify the minimum shaping rate for an adjusted queue. The queue is associated with a scheduler.

Options ***rate***—Minimum shaping rate for a node or a queue, in Mbps

$junos-cos-adjust-minimum—Junos OS predefined variable that is replaced with the minimum shaping rate for a node that is obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles]** hierarchy level.

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

Related Documentation

- Configuring a Dynamic Minimum Adjusted Shaping Rate on Scheduler Nodes on page 38
- Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 39

adjust-percent (Dynamic Schedulers)

Syntax adjust-percent *percentage*;

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

Release Information Statement introduced in Junos OS Release 11.4.

Description For a MPC/MIC interface, determine the percentage of adjustment for the shaping rate of a queue.

Options ***percentage***—Percentage of the shaping rate to adjust.

Range: 0 through 100 percent

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

Related Documentation

- Configuring a Dynamic Shaping-Rate Adjustment for Queues on page 39

ancp (Adjustment Control Profiles)

Syntax

```
ancp {
  priority priority;
  algorithm algorithm;
}
```

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

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the shaping rate adjustment controls for the ANCP application.

Options

priority—Priority of the ANCP application in the adjustment control profile.

Range: 1 through 10; 1 being the highest priority.

Default: 1

algorithm—Rate adjustment algorithm used by the ANCP application.

Values:

- adjust-never—Do not perform rate adjustments.
- adjust-always—Adjust the shaping rate unconditionally.
- adjust-less—Adjust the shaping rate if it is less than the configured value.
- adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
- adjust-greater—Adjust the shaping rate if it is greater than the configured value.
- adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Default: adjust-always

Required Privilege Level

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- adjustment-control-profiles on page 82
- application (Adjustment Control Profiles) on page 86

application (Adjustment Control Profiles)

Syntax

```
application {
  ancp;
  radius-coa;
  pppoe-tags;
}
```

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

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure which applications in the adjustment control profile can make shaping rate adjustments.

The remaining statements are explained separately.

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- adjustment-control-profiles on page 82

buffer-size (Dynamic Scheduling)

Syntax buffer-size (percent (*percentage* | $junos-cos-scheduler-bs) | remainder | temporal (*microseconds* | $junos-cos-scheduler-bs));

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

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-bs** predefined variable introduced in Junos OS Release 9.4.

Description Specify buffer size.

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.

Options

percent *percentage*—Buffer size as a percentage of total buffer.

remainder—Remaining buffer available.

temporal *microseconds*—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value.

Range: The ranges vary by platform as follows:

- For IQ PICs on M320 routers: 1 through 50,000 microseconds.
- For IQ PICs on other M Series routers: 1 through 100,000 microseconds.
- For other M Series routers: 1 through 200,000 microseconds.

$junos-scheduler-bs—Junos predefined variable that is replaced with the buffer size obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Schedulers in a Dynamic Profile for Subscriber Access*
- *scheduler (Dynamic Scheduler Maps)*

bytes (Dynamic Traffic Shaping)

Syntax bytes *bytes* | $junos-cos-byte-adjust;

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name* overhead-accounting],
[edit class-of-service traffic-control-profiles *profile-name* overhead-accounting]

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the number of overhead bytes.

Options ***bytes***—Byte adjustment value for the **cell-mode** or **frame-mode** shaping options. This can be the predefined variable **$junos-cos-byte-adjust**, which is the variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

BEST PRACTICE: We recommend using the cell-mode **cell-mode-bytes** ***cell-mode-bytes*** option or the frame-mode **frame-mode-bytes** ***frame-mode-bytes*** option rather than the **bytes** option.

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- **adjustment-control-profiles on page 82**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 51
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- *egress-shaping-overhead*
- **cell-mode on page 89**
- **frame-mode on page 98**

cell-mode (Dynamic Traffic Shaping)

Syntax cell-mode (bytes *bytes* | $junos-cos-byte-adjust | cell-mode-bytes *cell-mode-bytes* |$junos-cos-byte-adjust-cell);

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name* overhead-accounting],
[edit class-of-service traffic-control-profiles *profile-name* overhead-accounting],

Release Information Statement introduced in Junos OS Release 10.2.
Variable ***$junos-cos-byte-adjust-cell*** introduced in Junos OS Release 13.1.

Description Configure the mode to shape downstream ATM traffic as cells.

Options ***bytes***—Byte adjustment value for the **cell-mode** or **frame-mode** shaping options.

$junos-cos-byte-adjust—Predefined variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

cell-mode-bytes *cell-mode-bytes*—Shaping is based on the number of bytes in cells, and accounts for the ATM cell encapsulation and padding overhead. The resulting traffic stream conforms to the policing rates configured in downstream ATM switches, reducing the number of packet drops in the Ethernet network.

$junos-cos-byte-adjust-cell—Predefined variable for the cell mode shaping. This variable can not be used when the **overhead-accounting bytes *bytes*** option is configured.

BEST PRACTICE: We recommend using the **cell-mode-bytes *cell-mode-bytes*** option rather than the **bytes** option.

Range: –120 through 124 bytes

NOTE: If you specify a value for the **bytes *bytes*** option, you cannot specify a value for either the **cell-mode-bytes** option.

NOTE: Cell mode is supported only on logical interfaces and interface sets; it is not supported on physical interfaces (ifd or ifd-remaining).

Default: The default is frame-mode.

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- **adjustment-control-profiles on page 82**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 51
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- *egress-shaping-overhead*
- **bytes on page 88**
- **frame-mode on page 98**

class-of-service (Dynamic Profiles)

Syntax class-of-service { ... }

Hierarchy Level [edit dynamic-profiles *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure Junos OS CoS features in a dynamic profile.

Default If you do not configure any CoS features, all packets are transmitted from output transmission queue 0.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Static Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*
- *Configuring Dynamic Hierarchical Scheduling and Queuing in a Dynamic Profile for Subscriber Access*

delay-buffer-rate (Dynamic Traffic Shaping)

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

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

Release Information Statement introduced in Junos OS Release 9.2.
The **$junos-cos-delay-buffer-rate** variable introduced in Junos OS Release 9.4.

Description 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.

Options *rate*—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

$junos-cos-delay-buffer-rate—Junos predefined variable that is replaced with the delay-buffer rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- *output-traffic-control-profile*

dynamic-class-of-service-options (Dynamic Traffic Shaping)

Syntax

```
dynamic-class-of-service-options {
  vendor-specific-tags access-loop-encapsulation;
  vendor-specific-tags actual-data-rate-downstream;
}
```

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

Release Information Statement introduced in Junos OS Release 12.1.

Description Configure the shaping-rate and overhead-accounting class-of-service attributes based on access line parameters in PPPoE discovery packets on dynamic subscriber interfaces.

Options

vendor-specific-tags—Use Vendor-Specific Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101] to set the rate-shaping and overhead-accounting class-of-service attributes.

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

- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 17
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 46
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9

effective-shaping-rate

Syntax

```
effective-shaping-rate;
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 13.2.

Description Specify that the Cos-Effective-Shaping-Rate VSA [26–177] included in RADIUS Acct-Start, Acct-Stop, and Interim-Acct messages reports the actual rate of the downstream traffic for a subscriber, in kilobits per second.

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

Related Documentation

- Reporting the Effective Shaping Rate for Subscribers on page 46

excess-priority (Dynamic Schedulers)

Syntax excess-priority (low | high | $junos-cos-scheduler-excess-priority | none);

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

Release Information Statement introduced in Junos OS Release 10.2.
Option **none** introduced in Junos OS Release 11.4.

Description Determine the priority of excess bandwidth traffic on a scheduler in a dynamic profile.

Options

low—Excess traffic for this scheduler has low priority.

high—Excess traffic for this scheduler has high priority.

$junos-cos-scheduler-excess-priority—Variable for the excess-priority that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

none—System does not demote the priority of guaranteed traffic when the bandwidth exceeds the shaping rate or the guaranteed rate.

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

Related Documentation

- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 52
- *scheduler*

excess-rate (Dynamic Schedulers)

Syntax excess-rate percent (*percentage* | $junos-cos-scheduler-excess-rate);

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

Release Information Statement introduced in Junos OS Release 10.2.

Description Determine the percentage of excess bandwidth traffic to share.

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

Range: 0 through 100 percent

$junos-cos-scheduler-excess-rate—Variable for the excess rate that is specified for a scheduler. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 52
- *output-traffic-control-profile*

excess-rate (Dynamic Traffic Shaping)

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

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

Release Information Statement introduced in Junos OS Release 10.2.

Description For an MPC interface, determine the percentage or proportion of excess bandwidth traffic to share for all priorities of traffic.

Options

percentage—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

$junos-cos-excess-rate—Variable for the excess rate that is specified for the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 52
- *output-traffic-control-profile*

excess-rate-high (Dynamic Traffic Shaping)

Syntax excess-rate-high ((percent *percentage* | $junos-cos-excess-rate-high) | proportion *value*);

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

Release Information Statement introduced in Junos OS Release 11.4.

Description For an MPC/MIC interface, determine the percentage of excess bandwidth for high-priority traffic to share.

Options

percentage—Percentage of the excess bandwidth to share.

Range: 0 through 100 percent

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

$junos-cos-excess-rate-high—Variable for the excess rate that is specified for high-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 52
- *output-traffic-control-profile*

excess-rate-low (Dynamic Traffic Shaping)

Syntax excess-rate-low ((percent *percentage* | $junos-cos-excess-rate-low) | proportion *value*);

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

Release Information Statement introduced in Junos OS Release 11.4.

Description For an MPC/MIC interface, determine the percentage of excess bandwidth for low-priority traffic to share.

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

Range: 0 through 100 percent

value—Proportion of the excess bandwidth to share.

Range: 0 through 1000

$junos-cos-excess-rate-low—Variable for the excess rate that is specified for low-priority traffic on the logical interface. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- Managing Excess Bandwidth Distribution for Dynamic CoS on MIC and MIC Interfaces on page 52
- *output-traffic-control-profile*

frame-mode (Dynamic Traffic Shaping)

Syntax frame-mode (bytes | $junos-cos-byte-adjust | frame-mode-bytes *frame-mode-bytes* |$junos-cos-byte-adjust-frame);

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

Release Information Statement introduced in Junos OS Release 10.2.
Variable ***$junos-cos-byte-adjust-frame*** introduced in Junos OS Release 13.1.

Description Configure the mode to shape downstream ATM traffic based as frames.

Default The default is **frame-mode**.

Options ***bytes***—Byte adjustment value for the **cell-mode** or **frame-mode** shaping options.

$junos-cos-byte-adjust—Predefined variable for byte adjustment that is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

frame-mode-bytes *frame-mode-bytes*—Overhead bytes when in frame-mode. Traffic shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead.

$junos-cos-byte-adjust-frame—Predefined variable for frame mode shaping. This variable can not be used when the **overhead-accounting bytes *bytes*** option is configured.

BEST PRACTICE: We recommend using the **frame-mode-bytes *frame-mode-bytes*** option rather than the **bytes** option.

Range: –120 through 124 bytes

NOTE: If you specify a value for the **bytes *bytes*** option, you cannot specify a value for either the **frame-mode-bytes** option.

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- **adjustment-control-profiles on page 82**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 51

- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- *egress-shaping-overhead*
- bytes on page 88
- cell-mode on page 89

guaranteed-rate (Dynamic Traffic Shaping)

Syntax guaranteed-rate (*rate* | $junos-cos-guaranteed-rate) <burst-size [*bytes* | $junos-cos-guaranteed-rate-burst]>;

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

Release Information Statement introduced in Junos OS Release 9.2.
The **$junos-cos-guaranteed-rate** variable introduced in Junos OS Release 9.4.
Option **burst-size** introduced in Junos OS Release 11.4.

Description Configure a guaranteed minimum rate for a logical interface.

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 ***rate***—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

$junos-cos-guaranteed-rate—Junos predefined variable that is replaced with the guaranteed rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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

Range: 0 through 1,000,000,000

$junos-cos-guaranteed-rate-burst—(Optional) Variable for the burst-size that is specified for the guaranteed rate. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profile]** hierarchy level.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- *output-traffic-control-profile*

interfaces (Dynamic CoS Definition)

Syntax

```
interfaces {
  interface-name {
    unit logical-unit-number {
      classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        ieee-802.1 (classifier-name | default) vlan-tag (inner | outer)
        inet-precedence (classifier-name | default);
      }
      output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default);
      }
    }
  }
}
```

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

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure interface-specific CoS properties for incoming packets.

Options ***interface-name***—Either the specific name of the interface you want to assign to the dynamic profile or the interface variable ($junos-interface-ifd-name). The interface variable is dynamically replaced with the interface the client accesses when connecting to the router.

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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*

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 on page 57

multicast (Dynamic Routing Options)

Syntax

```
multicast {
  interface interface-name {
    no-qos-adjust;
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name* routing-options],
[edit dynamic-profiles *profile-name* routing-instances *routing-instance-name* routing-options]

NOTE: You cannot apply a scope policy to a specific routing instance. That is, all scoping policies are applied to all routing instances. However, the scope statement does apply individually to a specific routing instance.

Release Information Statement introduced in Junos OS Release 9.6.

Description Dynamically configure interface-specific multicast routing options properties.

The remaining statements are explained separately.

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

Related Documentation

- *Example: Configuring the Multicast Forwarding Cache*
- *Example: Configuring a Multicast Flow Map*
- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

no-qos-adjust (Dynamic Routing Options)

Syntax no-qos-adjust;

Hierarchy Level [edit dynamic-profiles *profile-name* routing-options multicast interface *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6.

Description Disable hierarchical bandwidth adjustment for all dynamically created subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

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

Related Documentation

- *Example: Configuring Multicast with Subscriber VLANs*

overhead-accounting (Dynamic Traffic Shaping)

Syntax

```
overhead-accounting {
  bytes bytes;
  cell-mode cell-mode-bytes cell-mode-bytes;
  frame-mode frame-mode-bytes frame-mode-bytes;
}
```

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

Release Information Statement introduced in Junos OS Release 10.2.

Description Configure the mode to shape downstream ATM traffic based on either frames or cells.

Default The default is frame-mode.

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

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- **adjustment-control-profiles on page 82**
- Configuring Dynamic Shaping Parameters to Account for Overhead in Downstream Traffic Rates on page 51
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9
- *egress-shaping-overhead*
- **bytes on page 88**
- **cell-mode on page 89**
- **frame-mode on page 98**

pppoe-tags (Adjustment Control Profiles)

Syntax

```
pppoe-tags {
  priority priority;
  algorithm algorithm;
}
```

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

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the shaping rate adjustment controls for the Point-to-Point Protocol over Ethernet (PPPoE) Tags application.

Options

priority—Priority of the Point to Point Protocol over Ethernet IA Tags application in the adjustment control profile.

Range: 1 through 10; 1 being the highest priority.

Default: 2

algorithm—Rate adjustment algorithm used by the Point to Point Protocol over Ethernet (PPPoE) IA Tags application.

Values:

- adjust-never—Do not perform rate adjustments.
- adjust-always—Adjust the shaping rate unconditionally.
- adjust-less—Adjust the shaping rate if it is less than the configured value.
- adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
- adjust-greater—Adjust the shaping rate if it is greater than the configured value.
- adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Default: adjust-less

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- adjustment-control-profiles on page 82
- application (Adjustment Control Profiles) on page 86

radius-coa (Adjustment Control Profiles)

Syntax

```
radius-coa {
  priority priority;
  algorithm algorithm;
}
```

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

Release Information Statement introduced in Junos OS Release 13.1.

Description Configure the shaping rate adjustment controls for the RADIUS CoA application.

Options

priority—Priority of the RADIUS CoA application in the adjustment control profile.

Range: 1 through 10; 1 being the highest priority.

Default: 1

algorithm—Rate adjustment algorithm used by the RADIUS CoA application.

Values:

- adjust-never—Do not perform rate adjustments.
- adjust-always—Adjust the shaping rate unconditionally.
- adjust-less—Adjust the shaping rate if it is less than the configured value.
- adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value.
- adjust-greater—Adjust the shaping rate if it is greater than the configured value.
- adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Default: adjust-always

Required Privilege Level

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

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- Verifying the CoS Adjustment Control Profile Configuration on page 119
- adjustment-control-profiles on page 82
- application (Adjustment Control Profiles) on page 86

routing-options (Dynamic Profiles)

Syntax

```
routing-options {
  access {
    route prefix {
      metric route-cost;
      next-hop next-hop;
      preference route-distance;
      tag route-tag;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
  multicast {
    interface interface-name {
      no-qos-adjust;
    }
  }
  rib routing-table-name {
    access {
      route prefix {
        metric route-cost;
        next-hop next-hop;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
  }
}
```

Hierarchy Level [edit dynamic-profiles *profile-name*],
[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]

Release Information Statement introduced in Junos OS Release 9.6.
Support at the **[edit dynamic-profiles *profile-name* routing-instances $junos-routing-instance]** hierarchy level introduced in Junos OS Release 10.1.

Description Configure protocol-independent routing properties in a dynamic profile.

The remaining statements are explained separately.

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

Related Documentation
- *Configuring Dynamic Access Routes for Subscriber Management*
- *Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management*

schedulers (Dynamic CoS Definition)

Syntax

```
schedulers {
  scheduler-name{
    adjust-minimum rate;
    adjust-percent percentage;
    buffer-size (percent percentage | remainder | temporal microseconds |
      $junos-cos-scheduler-bs);
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
      (any | non-tcp | tcp) drop-profile (profile-name | predefined-variable);
    excess-priority (low | high | $junos-cos-scheduler-excess-priority | none);
    excess-rate (percent percentage  | percent $junos-cos-scheduler-excess-rate);
    priority (priority-level | $junos-cos-scheduler-priority);
    shaping-rate (rate | predefined-variable) <burst-size  bytes>;
    transmit-rate (rate | percent percentage | remainder | percent percentage
      $junos-cos-scheduler-tx) <exact | rate-limit>;
  }
}
```

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

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler** predefined variable introduced in Junos OS Release 9.4.

Description Specify scheduler name and parameter values in a dynamic profile.

Options ***scheduler-name***—Name of the scheduler to be configured or the Junos OS predefined variable (**$junos-cos-scheduler**). The predefined variable is replaced with the scheduler name obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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
- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Schedulers in a Dynamic Profile for Subscriber Access*
- *scheduler*

shaping-rate (Dynamic Traffic Shaping and Scheduling)

Syntax shaping-rate (*rate* | *predefined-variable*) <burst-size *bytes* | $junos-cos-shaping-rate-burst>;

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

Release Information Statement introduced in Junos OS Release 9.2.
The **$junos-cos-shaping-rate** variable for traffic-control profiles introduced in Junos OS Release 9.4.
The **$junos-cos-scheduler-shaping-rate** variable for schedulers introduced in Junos OS Release 10.2.
Option **burst-size** introduced in Junos OS Release 11.4.

Description Configure a shaping rate for a logical interface or a scheduler. 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).

Options *rate*—Peak rate in bits per second (bps). You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

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

predefined-variable—One of the following Junos predefined variables. The variable is replaced with a value obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

- **$junos-cos-shaping-rate**—Variable for the shaping rate that is specified for the logical interface. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profiles *profile-name*]** hierarchy level.
- **$junos-cos-scheduler-shaping-rate**—Variable for the shaping rate that is specified for a scheduler. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service schedulers *scheduler-name*]** hierarchy level.

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

Range: 0 through 1,000,000,000

$junos-cos-shaping-rate-burst—(Optional) Variable for the burst-size that is specified for the shaping rate. Use this variable at the **[edit dynamic-profiles *profile-name* class-of-service traffic-control-profile]** hierarchy level.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- *output-traffic-control-profile*

targeted-distribution (Static Interfaces over Aggregated Ethernet)

Syntax targeted-distribution;

Hierarchy Level [edit interfaces demux0 unit *logical-unit-number*],
[edit interfaces pp0 unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 13.2R2 for EX Series switches.

Description Configure egress data for a logical interface to be sent across a single member link in an aggregated Ethernet bundle. A backup link is provisioned with CoS scheduling resources in the event that the primary assigned link goes down. The aggregated Ethernet interface must be configured without link protection.

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

Related Documentation

- *CoS for PPPoE Subscriber Interfaces Overview*
- *Configuring the Distribution Type for PPPoE Subscribers on Aggregated Ethernet Interfaces*
- *Verifying the Distribution of PPPoE Subscribers in an Aggregated Ethernet Interface*
- *Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis*
- *Configuring Module Redundancy for a Virtual Chassis*
- *Configuring Chassis Redundancy for a Virtual Chassis*

traffic-control-profiles (Dynamic CoS Definition)

Syntax

```
traffic-control-profiles profile-name {
  adjust-minimum rate;
  delay-buffer-rate (percent percentage | rate);
  excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
  excess-rate-high (percent percentage | proportion value);
  excess-rate-low (percent percentage | proportion value);
  guaranteed-rate (percent percentage | rate) <burst-size bytes>;
  overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
  scheduler-map map-name;
  shaping-rate (percent percentage | rate | predefined-variable) <burst-size bytes>;
}
```

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

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure traffic shaping and scheduling profiles.

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

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Traffic Scheduling and Shaping for Subscriber Access*
- *Using the CLI to Modify Traffic-Control Profiles That Are Currently Applied to Subscribers*
- *output-traffic-control-profile*

transmit-rate (Dynamic Schedulers)

Syntax transmit-rate (*rate* | percent *percentage* | remainder | percent *percentage* $junos-cos-scheduler-tx) <exact | rate-limit>;

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

Release Information Statement introduced in Junos OS Release 9.3.
The **$junos-cos-scheduler-tx** predefined variable introduced in Junos OS Release 9.4.

Description Specify the transmit rate or percentage for a scheduler in a dynamic profile.

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.

Options ***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

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

Range: 0 through 100 percent

remainder—Use remaining rate available.

$junos-cos-scheduler-tx—Junos predefined variable that is replaced with the transmission rate obtained from the RADIUS server when a subscriber authenticates over the interface to which the dynamic profile is attached.

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. Make sure this value never exceeds the rate-controlled amount.

rate-limit—(Optional) Limit the transmission rate to the rate-controlled amount during congestion. In contrast to the **exact** option, when there is no congestion, the scheduler with the **rate-limit** option shares unused bandwidth above the rate-controlled amount.

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Configuring Schedulers in a Dynamic Profile for Subscriber Access*
- *scheduler*

unit (Dynamic Traffic Shaping)

Syntax

```
unit logical-unit-number {
  classifiers {
    type (classifier-name | default);
  }
  output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
  rewrite-rules {
    dscp (rewrite-name | default);
    dscp-ipv6 (rewrite-name | default);
    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
    inet-precedence (rewrite-name | default);
    }
  }
}
```

Hierarchy Level

[edit dynamic-profiles *profile-name* class-of-service interfaces *interface-name*],
[edit dynamic-profiles *profile-name* interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.2.
Support at the **[edit dynamic-profiles *profile-name* class-of-service interfaces interface-set *interface-set-name*]** hierarchy level introduced in Junos OS Release 10.4.

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—One of the following options:

- **$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **$junos-interface-unit**—For dynamic demux and dynamic PPPoE interfaces, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP or PPP when it accesses the subscriber network.
- *value*—Specific unit number of the interface you want to assign to the dynamic-profile

Range: 0 through 16385. For demux and PPPoE interfaces, the unit numbers can range from 0 through 1,073,741,823.

The remaining statements are explained separately. The **classifiers**, **output-traffic-control-profile**, and **rewrite-rules** statements are not supported for interface sets.

Required Privilege Level

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

Related Documentation

- Guidelines for Configuring Dynamic CoS for Subscriber Access on page 27
- *Applying Traffic Shaping and Scheduling to a Subscriber Interface in a Dynamic Profile*
- *Configuring an Interface Set of Subscribers in a Dynamic Profile*

vendor-specific-tags (Dynamic Traffic Shaping)

Syntax

```
vendor-specific-tags actual-data-rate-downstream;
vendor-specific-tags access-loop-encapsulation;
```

Hierarchy Level [edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]

Release Information Statement introduced in Junos OS Release 12.1.

Description Set the shaping-rate and overhead-accounting class-of-service attributes based on Vendor-Specific Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101].

Options **vendor-specific-tags** can be set to one or both of the following:

- **access-loop-encapsulation**—Set the overhead-accounting class-of-service attribute based on access line parameters in PPPoE discovery packets on dynamic subscriber interfaces.
- **actual-data-rate-downstream**—Set the shaping-rate class-of-service attribute based on the actual-data-rate-downstream attribute.

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

Related Documentation

- Setting Class-of-Service Parameters Using PPPoE Vendor-Specific Tags on page 17
- Configuring the Shaping Rate and Overhead Accounting Based on PPPoE Vendor-Specific Tags on Dynamic Subscriber Interfaces on page 46
- Bandwidth Management for Downstream Traffic in Edge Networks Overview on page 9

PART 3

Administration

CHAPTER 11

Monitoring CoS for Subscriber Access

Verifying the Scheduling and Shaping Configuration for Subscriber Access

Purpose View the class-of-service (CoS) configurations that are referenced in a dynamic profile for subscriber access.

Action

- To display the entire CoS configuration, including static and dynamic parameters:

  user@host> show class-of-service

- To display the CoS configuration for a subscriber interface:

  user@host> show class-of-service interface

- To display traffic shaping and scheduling profiles:

  user@host> show class-of-service traffic-control-profile

- To display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry:

  user@host> show class-of-service scheduler-map

Verifying the Number of Dedicated Queues Configured on MIC and MPC Interfaces

Purpose Display the number of dedicated queue resources that are configured for the logical interfaces on a port.

Action

```
user@host#show class-of-service interface ge-1/1/0
Physical interface: ge-1/1/0, Index: 166
Queues supported: 4, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2
  Chassis scheduler map: <default-chassis>, Index: 4

  Logical interface: ge-1/1/0.100, Index: 72, Dedicated Queues: no
    Shaping rate: 32000
    Object             Name                   Type                  Index
    Scheduler-map      <remaining>                                  0
    Classifier         ipprec-compatibility   ip                    13

  Logical interface: ge-1/1/0.101, Index: 73, Dedicated Queues: no
    Shaping rate: 32000
    Object             Name                   Type                  Index
    Scheduler-map      <remaining>                                  0
    Classifier         ipprec-compatibility   ip                    13

  Logical interface: ge-1/1/0.102, Index: 74, Dedicated Queues: yes
    Shaping rate: 32000
    Object                  Name               Type                 Index
    Traffic-control-profile <control_tc_prof>  Output               45866
```

Related Documentation

- *Managing Dedicated and Remaining Queues for Static CoS Configurations on MIC and MPC Interfaces*
- Managing Dedicated and Remaining Queues for Dynamic CoS Configurations on MIC and MPC Interfaces on page 57

Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface

Purpose View the distribution status of subscribers that are targeted to links in an aggregated Ethernet interface.

Action

- To display a summary of the distribution of links on the demux interface:

  user@host> **show interfaces demux0 extensive**

- To display the targeted distribution on a specific aggregated Ethernet interface:

  user@host> **show interfaces targeting ae*x***

Related Documentation

- *Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces*
- *Configuring Rebalancing of Demux Subscribers in an Aggregated Ethernet Interface*

Verifying the Configuration of Shaping-Rate Adjustments for Subscriber Local Loops

Purpose Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical CoS.

NOTE: After shaping-rate adjustments are enabled and the router has performed shaping-rate adjustments on a scheduler node, you can configure a new shaping rate by including the **shaping-rate** statement in a traffic-control profile and then applying that profile to that scheduler node. However, this new shaping-rate value does not immediately result in shaping traffic at the new rate. The scheduler node continues to be shaped at rate set by ANCP. Only when the ANCP shaping-rate adjustment feature is disabled is the scheduler node shaped at the newly configured shaping-rate.

Action Issue the **show class-of-service interface-set** operational command.

Related Documentation
- Enabling Shaping-Rate Adjustments for Subscriber Local Loops on page 40

Verifying the Configuration of ANCP for Shaping-Rate Adjustments

Purpose Use to display or clear information about the ANCP configuration for shaping-rate adjustments.

Action
- To display ANCP neighbor information, issue the ***show ancp neighbor*** operational command.
- To clear ANCP neighbors, issue the ***clear ancp neighbor*** operational command.
- To display ANCP subscriber information, issue the ***show ancp subscriber*** operational command.
- To display ANCP class-of-service information, issue the ***show ancp cos*** operational command.

If ANCP is not yet enabled, the process starts when you commit a configuration that contains the **protocols ancp** stanza.

Related Documentation
- *ANCP and the ANCP Agent Overview*
- *Configuring the ANCP Agent*

Verifying the CoS Adjustment Control Profile Configuration

Purpose View the class-of-service (CoS) adjustment control profile.

Action
- To display the CoS adjustment control profile:

  user@host> **show class-of-service adjustment-control-profile** *profile-name*

```
user@host> show class-of-service adjustment-control-profile acp1
  name: ANCP, priority: 1, algorithm: less
  name: RADIUS CoA, priority: 1, algorithm: always
  name: PPPoE IA tags, priority: 2, algorithm: less
```

```
user@host>
```

Related Documentation

- CoS Adjustment Control Profiles Overview on page 6
- Configuring CoS Adjustment Control Profiles on page 35
- adjustment-control-profiles on page 82
- application (Adjustment Control Profiles) on page 86

Verifying the Effective Shaping Rate Reporting Configuration

Purpose Verify whether reporting is enabled for the effective shaping rate. Display the effective shaping rate when reporting is enabled.

Action

- To display configuration information for effective shaping rate reporting:

  ```
  [edit]
  user@host# show chassis
  ...
  effective-shaping-rate;
  ...
  ```

- To display the effective shaping rate in kilobits per second when reporting is enabled:

  ```
  user@host> show subscribers extensive
  Type: VLAN
  Logical System: default
  Routing Instance: default
  Interface: demux0.1073741837
  Interface type: Dynamic
  Interface Set: ifset-1
  Underlying Interface: ae1
  Dynamic Profile Name: svlan-dhcp-test
  State: Active
  Session ID: 1
  Stacked VLAN Id: 0x8100.201
  VLAN Id: 0x8100.201
  Login Time: 2011-11-30 00:18:04 PST
  Effective shaping-rate: 31000000
  ...
  ```

Related Documentation

- Reporting the Effective Shaping Rate for Subscribers on page 46

CHAPTER 12

Monitoring Commands

show class-of-service

Syntax show class-of-service

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display the entire class-of-service (CoS) configuration, including system-chosen defaults. Executing this command is equivalent to executing all **show class-of-service** commands in succession.

Options This command has no options.

Required Privilege Level view

List of Sample Output show class-of-service on page 122

Output Fields See the output field descriptions for the commands.

Sample Output

show class-of-service

```
user@host> show class-of-service
Forwarding class                      Queue
  best-effort                             0
  expedited-forwarding                    1
  assured-forwarding                      2
  network-control                         3
Code point type: dscp
  Alias              Bit pattern
  af11               001010
  af12               001100
  af13               001110
...
Code point type: dscp-ipv6
  Alias              Bit pattern
  af11               001010
  af12               001100
  af13               001110
...
Code point type: exp
  Alias              Bit pattern
  af11               100
  af12               101
  be                 000
...
Code point type: ieee-802.1
  Alias              Bit pattern
  af11               100
  af12               101
  be                 000
...
Classifier: dscp-default, Code point type: dscp, Index: 6
  Code point         Forwarding class                  Loss priority
  000000             best-effort                       low
```

```
  000001             best-effort                        low
  000010             best-effort                        low
....
Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 7
  Code point         Forwarding class                   Loss priority
  000000             best-effort                        low
  000001             best-effort                        low
  000010             best-effort                        low
...
Loss-priority-map: frame-relay-de-default, Code point type: frame-relay-de, Index:
 12
  Code point         Loss priority
  0                  low
  1                  high

Rewrite rule: dscp-default, Code point type: dscp, Index: 23
  Forwarding class                   Loss priority       Code point
  best-effort                        low                 000000
  best-effort                        high                000000
  expedited-forwarding               low                 101110
...
Rewrite rule: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 24
  Forwarding class                   Loss priority       Code point
  best-effort                        low                 000000
  best-effort                        high                000000
...
....
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100

Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: best-effort, Index: 16
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent, Priority:
 low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    <default-drop-profile>
      Medium low      any             1    <default-drop-profile>
      Medium high     any             1    <default-drop-profile>
      High            any             1    <default-drop-profile>
...
Physical interface: fe-0/0/0, Index: 137
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2

  Logical interface: fe-0/0/0.0, Index: 69
    Object            Name                   Type                    Index
    Adaptive-shaper   fr-shaper                                      35320
    Classifier        ipprec-compatibility   ip                         11

Physical interface: fe-0/0/1, Index: 138
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
...
```

show class-of-service interface

Syntax

show class-of-service interface
<**comprehensive** | **detail**> <*interface-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Forwarding class map information added in Junos OS Release 9.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.
Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Options **detail** and **comprehensive** introduced in Junos OS Release 11.4.

Description

Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

Options

none—Display CoS associations for all physical and logical interfaces.

comprehensive—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.

detail—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.

If the **interface** *interface-name* is a physical interface, the output includes:

- Brief QoS information about the physical interface
- Brief QoS information about the logical interface
- CoS information about the physical interface
- Brief information about filters or policers of the logical interface
- Brief CoS information about the logical interface

If the **interface** *interface-name* is a logical interface, the output includes:

- Brief QoS information about the logical interface
- Information about filters or policers for the logical interface
- CoS information about the logical interface

interface-name—(Optional) Display class-of-service (CoS) associations for the specified interface.

Required Privilege Level

view

List of Sample Output

show class-of-service interface (Physical) on page 135

show class-of-service interface (Logical) on page 136
show class-of-service interface (Gigabit Ethernet) on page 136
show class-of-service interface (PPPoE Interface) on page 136
show class-of-service interface (T4000 Routers with Type 5 FPCs) on page 136
show class-of-service interface detail on page 137
show class-of-service interface comprehensive on page 137
show class-of-service interface (ACX Series Routers) on page 148

Output Fields Table 13 on page 125 describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

Table 13: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	Index of this interface or the internal index of this object.
Dedicated Queues	Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Total non-default queues created	Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.
Rewrite Input IEEE Code-point	(QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the **Shaping rate** field is displayed for either the physical interface or the logical interface.
Scheduler map	Name of the output scheduler map associated with this interface.
Scheduler map forwarding class sets	(QFX Series only) Name of the fabric forwarding class set scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Classifier	Name and type of classifiers associated with this interface.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	(QFX Series only) Congestion notification state, **enabled** or **disabled**.
Logical interface	Name of a logical interface.
Object	Category of an object: **Classifier**, **Fragmentation-map** (for LSQ interfaces only), **Scheduler-map**, **Rewrite**, or **Translation Table** (for IQE PICs only).
Name	Name of an object.
Type	Type of an object: **dscp**, **dscp-ipv6**, **exp**, **ieee-802.1**, **ip**, or **inet-precedence**.
Link-level type	Encapsulation on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status. • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Device flags	The **Device flags** field provides information about the physical device and displays one or more of the following values: • **Down**—Device has been administratively disabled. • **Hear-Own-Xmit**—Device receives its own transmissions. • **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint. • **Loopback**—Device is in physical loopback. • **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback. • **No-Carrier**—On media that support carrier recognition, no carrier is currently detected. • **No-Multicast**—Device does not support multicast traffic. • **Present**—Device is physically present and recognized. • **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • **Quench**—Transmission on the device is quenched because the output buffer is overflowing. • **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • **Running**—Device is active and enabled.

Interface flags	The **Interface flags** field provides information about the physical interface and displays one or more of the following values:
	• **Admin-Test**—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
	• **Disabled**—Interface is administratively disabled.
	• **Down**—A hardware failure has occurred.
	• **Hardware-Down**—Interface is nonfunctional or incorrectly connected.
	• **Link-Layer-Down**—Interface keepalives have indicated that the link is incomplete.
	• **No-Multicast**—Interface does not support multicast traffic.
	• **No-receive No-transmit**—Passive monitor mode is configured on the interface.
	• **Point-To-Point**—Interface is point-to-point.
	• **Pop all MPLS labels from packets of depth**—MPLS labels are removed as packets arrive on an interface that has the **pop-all-labels** statement configured. The depth value can be one of the following:
	  ◦ **1**—Takes effect for incoming packets with one label only.
	  ◦ **2**—Takes effect for incoming packets with two labels only.
	  ◦ **[1 2]**—Takes effect for incoming packets with either one or two labels.
	• **Promiscuous**—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
	• **Recv-All-Multicasts**—Interface is in multicast promiscuous mode and provides no multicast filtering.
	• **SNMP-Traps**—SNMP trap notifications are enabled.
	• **Up**—Interface is enabled and operational.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Flags	The **Logical interface flags** field provides information about the logical interface and displays one or more of the following values: • **ACFC Encapsulation**—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • **Device-down**—Device has been administratively disabled. • **Disabled**—Interface is administratively disabled. • **Down**—A hardware failure has occurred. • **Clear-DF-Bit**—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • **Hardware-Down**—Interface protocol initialization failed to complete successfully. • **PFC**—Protocol field compression is enabled for the PPP session. • **Point-To-Point**—Interface is point-to-point. • **SNMP-Traps**—SNMP trap notifications are enabled. • **Up**—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (**Up** or **Down**)
Link	Status of physical link (**Up** or **Down**).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.
Link flags	Provides information about the physical link and displays one or more of the following values: • **ACFC**—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • **Give-Up**—Link protocol does not continue connection attempts after repeated failures. • **Loose-LCP**—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • **Loose-LMI**—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • **Loose-NCP**—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • **Keepalives**—Link protocol keepalives are enabled. • **No-Keepalives**—Link protocol keepalives are disabled. • **PFC**—Protocol field compression is configured. The PPP session negotiates the PFC option.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
CoS queues	Number of CoS queues configured.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: *year-month-day hour:minute:second:timezone* (*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**
Statistics last cleared	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.
Input errors	Input errors on the interface. The labels are explained in the following list: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Bucket Drops**—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **HS link FIFO overflows**—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Output errors	Output errors on the interface. The labels are explained in the following list: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Drops** field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **HS link FIFO underflows**—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeds the MTU of the interface.
Egress queues	Total number of egress queues supported on the specified interface.
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism. NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the **Dropped packets** field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.
SONET alarms **SONET defects**	(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: **SONET PHY**, **SONET section**, **SONET line**, and **SONET path**.
SONET PHY	Counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. A state other than **OK** indicates a problem. The **SONET PHY** field has the following subfields: • **PLL Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
SONET section	Counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. A state other than **OK** indicates a problem. The **SONET section** field has the following subfields: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOS**—Loss of signal • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • **SES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)
SONET line	Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. A state other than **OK** indicates a problem. The **SONET line** field has the following subfields: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
SONET path	Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. A state other than **OK** indicates a problem. The **SONET path** field has the following subfields: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload (signal) label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **ES-PFE**—Errored seconds (far-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)
Received SONET overhead **Transmitted SONET overhead**	Values of the received and transmitted SONET overhead: • **C2**—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • **F1**—Section user channel byte. This byte is set aside for the purposes of users. • **K1** and **K2**—These bytes are allocated for APS signaling for the protection of the multiplex section. • **J0**—Section trace. This byte is defined for STS-1 number 1 of an STS-*N* signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • **S1**—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-*N* signal. • **Z3** and **Z4**—Allocated for future use.
Received path trace **Transmitted path trace**	SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.
HDLC configuration	Information about the HDLC configuration. • **Policing bucket**—Configured state of the receiving policer. • **Shaping bucket**—Configured state of the transmitting shaper. • **Giant threshold**—Giant threshold programmed into the hardware. • **Runt threshold**—Runt threshold programmed into the hardware.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. • **PLP byte**—Packet Level Protocol byte.
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
Forwarding classes	Total number of forwarding classes supported on the specified interface.
Egress queues	Total number of egress queues supported on the specified interface.
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
Tail-dropped packets	Number of packets dropped because of tail drop.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description

RED-dropped packets	Number of packets dropped because of random early detection (RED).
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories:
	  • **Low, non-TCP**—Number of low-loss priority non-TCP packets dropped because of RED.
	  • **Low, TCP**—Number of low-loss priority TCP packets dropped because of RED.
	  • **High, non-TCP**—Number of high-loss priority non-TCP packets dropped because of RED.
	  • **High, TCP**—Number of high-loss priority TCP packets dropped because of RED.
	• (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories:
	  • **Low**—Number of low-loss priority packets dropped because of RED.
	  • **Medium-low**—Number of medium-low loss priority packets dropped because of RED.
	  • **Medium-high**—Number of medium-high loss priority packets dropped because of RED.
	  • **High**—Number of high-loss priority packets dropped because of RED.
	NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.

RED-dropped bytes	Number of bytes dropped because of RED. The byte counts vary by PIC type.
	• (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories:
	  • **Low, non-TCP**—Number of low-loss priority non-TCP bytes dropped because of RED.
	  • **Low, TCP**—Number of low-loss priority TCP bytes dropped because of RED.
	  • **High, non-TCP**—Number of high-loss priority non-TCP bytes dropped because of RED.
	  • **High, TCP**—Number of high-loss priority TCP bytes dropped because of RED.
	NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.

Transmit rate	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
Rate Limit	Rate limiting configuration of the queue. Possible values are : • **None**—No rate limit. • **exact**—Queue transmits at the configured rate.
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as **low** or **high**.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: **low**, **medium-low**, **medium-high**, **high**, or **none**.

Table 13: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Drop profiles	Display the assignment of drop profiles. • **Loss priority**—Packet loss priority for drop profile assignment. • **Protocol**—Transport protocol for drop profile assignment. • **Index**—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • **Name**—Name of the drop profile. • **Type**—Type of the drop profile: **discrete** or **interpolated**. • **Fill Level**—Percentage fullness of a queue. • **Drop probability**—Drop probability at this fill level.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	Display the assignment of drop profiles. • **Loss priority**—Packet loss priority for drop profile assignment. • **Protocol**—Transport protocol for drop profile assignment. • **Index**—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • **Name**—Name of the drop profile. • **Type**—Type of the drop profile: **discrete** or **interpolated**. • **Fill Level**—Percentage fullness of a queue. • **Drop probability**—Drop probability at this fill level.
Adjustment information	Display the assignment of shaping-rate adjustments on a scheduler node or queue. • **Adjusting application**—Application that is performing the shaping-rate adjustment. ◦ The adjusting application can appear as **ancp LS-0**, which is the Junos OS Access Node Control Profile process (**ancpd**) that performs shaping-rate adjustments on schedule nodes. ◦ The adjusting application can also appear as **pppoe**, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode). • **Adjustment type**—Type of adjustment: **absolute** or **delta**. • **Configured shaping rate**—Shaping rate configured for the scheduler node or queue. • **Adjustment value**—Value of adjusted shaping rate. • **Adjustment target**—Level of shaping-rate adjustment performed: **node** or **queue**. • **Adjustment overhead-accounting mode**—Configured shaping mode: **frame** or **cell**.

Sample Output

show class-of-service interface (Physical)

```
user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4
```

```
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2032638653

  Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                  Name                  Type           Index
    Scheduler-map           <default>                             27
    Rewrite                 exp-default           exp             21
    Classifier              exp-default           exp             5
    Classifier              ipprec-compatibility  ip              8
    Forwarding-class-map    exp-default           exp             5
```

show class-of-service interface (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                  Name                  Type             Index
    Scheduler-map           <default>                               27
    Rewrite                 exp-default           exp               21
    Classifier              exp-default           exp               5
    Classifier              ipprec-compatibility  ip                8
    Forwarding-class-map    exp-default           exp               5
```

show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Input scheduler map: <default>, Index: 3
  Chassis scheduler map: <default-chassis>, Index: 4
```

show class-of-service interface (PPPoE Interface)

```
user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85
    Object                  Name                  Type           Index
    Traffic-control-profile tcp-pppoe.o.pp0.1     Output         2726446535
    Classifier              ipprec-compatibility  ip             13

    Adjusting application: PPPoE
      Adjustment type: absolute
      Adjustment value: 5000000
      Adjustment overhead-accounting mode: cell
      Adjustment target: node
```

show class-of-service interface (T4000 Routers with Type 5 FPCs)

```
user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
    Queues supported: 8, Queues in use: 4
      Shaping rate: 5000000000 bps
      Scheduler map: <default>, Index: 2
      Congestion-notification: Disabled

      Logical interface: xe-4/0/0.0, Index: 77
        Object                  Name                    Type
Index
        Classifier              ipprec-compatibility    ip
  13
```

show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Physical interface: ge-0/3/0, Index: 138
  Queues supported: 4, Queues in use: 5
  Shaping rate: 50000 bps
  Scheduler map: interface-schedular-map, Index: 58414
  Input shaping rate: 10000 bps
878674  Input scheduler map: schedular-map, Index: 15103
  Chassis scheduler map: <default-chassis>, Index: 4
  Congestion-notification: Disabled

  Logical interface ge-0/3/0.0
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    inet
    mpls
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.0      up    up   inet
                           mpls
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.0      up    up   inet
                           mpls

  Logical interface: ge-0/3/0.0, Index: 68
    Object                  Name                    Type               Index
    Rewrite                 exp-default             exp (mpls-any)       33
    Classifier              exp-default             exp                  10
    Classifier              ipprec-compatibility    ip                   13


  Logical interface ge-0/3/0.1
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    inet
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.1      up    up   inet
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.1      up    up   inet

  Logical interface: ge-0/3/0.1, Index: 69
    Object                  Name                    Type            Index
    Classifier              ipprec-compatibility    ip               13
```

show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601, Generation: 141
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```
  CoS queues     : 4 supported, 4 maximum usable queues
  Schedulers     : 256
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
  Last flapped   : 2010-09-07 06:35:22 PDT (15:14:42 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 total statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Drop   bytes  :                   0                   0 bps
   Drop   packets:                   0                   0 pps
  Label-switched interface (LSI) traffic statistics:
   Input  bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 af3                            0                   0                   0

    1 af2                            0                   0                   0

    2 ef2                            0                   0                   0

    3 ef1                            0                   0                   0

  Egress queues: 4 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 af3                            0                   0                   0

    1 af2                            0                   0                   0

    2 ef2                            0                   0                   0

    3 ef1                            0                   0                   0

  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
```

```
    CRC/Align errors                          0                0
    FIFO errors                               0                0
    MAC control frames                        0                0
    MAC pause frames                          0                0
    Oversized frames                          0
    Jabber frames                             0
    Fragment frames                           0
    VLAN tagged frames                        0
    Code violations                           0
  Filter statistics:
    Input packet count                        0
    Input packet rejects                      0
    Input DA rejects                          0
    Input SA rejects                          0
    Output packet count                                        0
    Output packet pad count                                    0
    Output packet error count                                  0
    CAM destination filters: 0, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
       Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
 OK
    Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 0
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %           usec
    2 ef2                    39          19500     0            120     high
none
    Direction : Input
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                              %            bps     %           usec
    0 af3                    30           3000    45              0      low
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                      0                      0 pps
    Bytes                :                      0                      0 bps
  Transmitted:
    Packets              :                      0                       0 pps
    Bytes                :                      0                      0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                      0                      0 pps
    RED-dropped bytes    :                      0                      0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                      0                      0 pps
    Bytes                :                      0                      0 bps
  Transmitted:
    Packets              :                      0                       0 pps
```

```
    Bytes               :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
  Transmitted:
    Packets             :                    0                    0 pps
    Bytes               :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
```

```
    RED-dropped bytes    :                    0                    0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 1, Forwarding classes: af2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets              :               108546                    0 pps
    Bytes                :             12754752                  376 bps
  Transmitted:
    Packets              :               108546                    0 pps
    Bytes                :             12754752                  376 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  : Not Available
    RED-dropped bytes    : Not Available

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
  Shaping rate: 50000 bps
```

```
Scheduler map: interface-schedular-map, Index: 58414

  Scheduler: ef2, Forwarding class: ef2, Index: 39155
    Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer
Limit: none, Priority: high
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
  Input shaping rate: 10000 bps
  Input scheduler map: schedular-map

Scheduler map: schedular-map, Index: 15103

  Scheduler: af3, Forwarding class: af3, Index: 35058
    Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer
 Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         40582    green
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any         18928    yellow
Drop profile: green, Type: discrete, Index: 40582
  Fill level    Drop probability
         50                   0
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: yellow, Type: discrete, Index: 18928
  Fill level    Drop probability
         50                   0
        100                 100
  Chassis scheduler map: < default-drop-profile>
Scheduler map: < default-drop-profile>, Index: 4

  Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
```

```
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100

  Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100

  Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
```

```
  Fill level    Drop probability
        100                 100

  Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
    Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
 Limit: none, Priority: low
    Excess Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any             1    < default-drop-profile>
      Medium low      any             1    < default-drop-profile>
      Medium high     any             1    < default-drop-profile>
      High            any             1    < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
  Congestion-notification: Disabled
Forwarding class                        ID      Queue  Restricted queue  Fabric
priority  Policing priority
  af3                                    0       0          0             low
            normal
  af2                                    1       1          1             low
            normal
  ef2                                    2       2          2             high
            normal
  ef1                                    3       3          3             high
            normal
  af1                                    4       4          0             low
            normal

  Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol inet, MTU: 1500, Generation: 172, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Input Filters: filter-in-ge-0/3/0.0-i,
      Policer: Input: p1-ge-0/3/0.0-inet-i
    Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0
```

```
      Flags: Is-Primary
      Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class ID  Queue  Restricted queue  Fabric priority  Policing priority
   SPU priority
best-effort      0   0      0                 low              normal
    low

Aggregate Forwarding-class statistics per forwarding-class
  Aggregate Forwarding-class statistics:
   Forwarding-class statistics:

Forwarding-class best-effort statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

IPv4 protocol forwarding-class statistics:
  Forwarding-class statistics:
  Forwarding-class best-effort statistics:

    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0
```

```
    IPv6  protocol forwarding-class statistics:
    Forwarding-class statistics:
     Forwarding-class best-effort statistics:

    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:
    Input unicast bytes:      0
    Output unicast bytes:     0
    Input unicast packets:    0
    Output unicast packets:   0

    Input multicast bytes:    0
    Output multicast bytes:   0
    Input multicast packets:  0
    Output multicast packets: 0


 Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0

Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.0      up    up   inet  filter-in-ge-0/3/0.0-i
                           mpls                       exp-filter
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.0      up    up
                           inet  p1-ge-0/3/0.0-inet-i
                           mpls

Filter: filter-in-ge-0/3/0.0-i
Counters:
Name                                                Bytes              Packets
count-filter-in-ge-0/3/0.0-i                            0                    0

Filter: exp-filter
Counters:
Name                                                Bytes              Packets
count-exp-seven-match                                   0                    0
count-exp-zero-match                                    0                    0
Policers:
Name                                               Packets
p1-ge-0/3/0.0-inet-i                                    0


  Logical interface: ge-0/3/0.0, Index: 68
    Object                 Name                   Type                    Index

    Rewrite                exp-default            exp (mpls-any)            33

Rewrite rule: exp-default, Code point type: exp, Index: 33
```

```
  Forwarding class                  Loss priority       Code point
  af3                               low                 000
  af3                               high                001
  af2                               low                 010
  af2                               high                011
  ef2                               low                 100
  ef2                               high                101
  ef1                               low                 110
  ef1                               high                111
    Object                 Name                  Type                      Index

    Classifier             exp-default           exp                         10

Classifier: exp-default, Code point type: exp, Index: 10
  Code point         Forwarding class                  Loss priority
  000                af3                               low
  001                af3                               high
  010                af2                               low
  011                af2                               high
  100                ef2                               low
  101                ef2                               high
  110                ef1                               low
  111                ef1                               high
    Object                 Name                  Type                      Index

    Classifier             ipprec-compatibility  ip                          13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                  Loss priority
  000                af3                               low
  001                af3                               high
  010                af3                               low
  011                af3                               high
  100                af3                               low
  101                af3                               high
  110                ef1                               low
  111                ef1                               high
Forwarding class                    ID      Queue  Restricted queue  Fabric
priority  Policing priority
  af3                               0       0          0             low
            normal
  af2                               1       1          1             low
            normal
  ef2                               2       2          2             high
            normal
  ef1                               3       3          3             high
            normal
  af1                               4       4          0             low
            normal

  Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Local statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
```

```
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 1500, Generation: 174, Route table: 0
      Flags: Sendbcast-pkt-to-re

  Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0


Interface       Admin Link Proto Input Filter         Output Filter
ge-0/3/0.1      up    up   mpls
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/3/0.1      up    up
                           mpls

  Logical interface: ge-0/3/0.1, Index: 69
    Object                 Name                   Type                    Index

    Classifier             ipprec-compatibility   ip                        13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
  Code point         Forwarding class                    Loss priority
  000                af3                                 low
  001                af3                                 high
  010                af3                                 low
  011                af3                                 high
  100                af3                                 low
  101                af3                                 high
  110                ef1                                 low
  111                ef1                                 high
Forwarding class                      ID      Queue  Restricted queue  Fabric
priority  Policing priority
  af3                                  0       0          0             low
            normal
  af2                                  1       1          1             low
            normal
  ef2                                  2       2          2             high
            normal
  ef1                                  3       3          3             high
            normal
  af1                                  4       4          0             low
            normal
```

show class-of-service interface (ACX Series Routers)

```
user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: at-0/0/0.0, Index: 69
```

```
  Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133
Queues supported: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled

  Logical interface: at-0/0/1.0, Index: 71

  Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Rewrite                 dscp-default            dscp                       31
Classifier              d1                      dscp                    11331
Classifier              ci                      ieee8021p                 583

  Logical interface: ge-0/1/0.0, Index: 73
Object                  Name                    Type                    Index
Rewrite                 custom-exp              exp (mpls-any)          46413

  Logical interface: ge-0/1/0.1, Index: 74

  Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Rewrite                 ri                      ieee8021p (outer)       35392
Classifier              ci                      ieee8021p                 583

Physical interface: ge-0/1/3, Index: 149
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
Classifier              ipprec-compatibility    ip                         13

  Logical interface: ge-0/1/3.0, Index: 77
Object                  Name                    Type                    Index
Rewrite                 custom-exp2             exp (mpls-any)          53581

Physical interface: ge-0/1/4, Index: 150
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                    Type                    Index
```

```
Classifier              ipprec-compatibility   ip                       13

Physical interface: ge-0/1/5, Index: 151
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              ipprec-compatibility   ip                       13

Physical interface: ge-0/1/6, Index: 152
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              ipprec-compatibility   ip                       13

Physical interface: ge-0/1/7, Index: 153
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              d1                     dscp                     11331

Physical interface: ge-0/2/0, Index: 154
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              ipprec-compatibility   ip                       13

Physical interface: ge-0/2/1, Index: 155
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              ipprec-compatibility   ip                       13

  Logical interface: ge-0/2/1.0, Index: 78

  Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              ipprec-compatibility   ip                       13

  Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object                  Name                   Type                     Index
Classifier              ipprec-compatibility   ip                       13

  Logical interface: xe-0/3/1.0, Index: 81

[edit]
user@host-g11#
```

show class-of-service interface-set

Syntax **show class-of-service interface-set**
<interface-set-name>

Release Information Command introduced in Junos OS Release 9.4.

Description Display the configured shaping rate and the adjusted shaping rate for each logical interface set configured for hierarchical class of service (CoS).

Options **none**—Display CoS associations for all logical interface sets.

interface-set *interface-set-name*—(Optional) Display CoS associations for the specified interface set.

Required Privilege Level view

List of Sample Output show class-of-service interface-set on page 153

Output Fields Table 14 on page 152 describes the output fields for the **show class-of-service interface-set** command. Output fields are listed in the approximate order in which they appear.

Table 14: show class-of-service interface-set Output Fields

Field Name	Field Description
Interface-set	Name of a logical interface set composed of one or more logical interfaces for which hierarchical scheduling is enabled.
Index	Index number of this interface set or the internal index number of this object.
Physical interface	Name of a physical interface.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Output traffic control profile	Name of the output traffic-control profile attached to the logical interface set.

Table 14: show class-of-service interface-set Output Fields *(continued)*

Field Name	Field Description
Adjusting application	Name of the application that communicates shaping-rate adjustment information to the Junos OS class-of-service process (**cosd**) on the broadband services router (BSR). The BSR uses the information from this application to perform shaping-rate adjustments on the scheduler node that manages the interface set. The adjusting application appears as **ancp LS-0** which is the Junos OS Access Node Control Profile process (**ancpd**) that performs shaping-rate adjustments on schedule nodes. The nodes are logical interface sets configured to represent subscriber local loops. When the synchronization speed of the DSL line changes, **ancpd** communicates the local loop speed to **cosd** over the default logical system, **LS-0**, and then the BSR throttles the shaping rate on the scheduler node to the loop speed. The adjusting application can also appear as **PPPoE**, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual data rate downstream attribute. The overhead accounting value is based on the access loop encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).
Adjustment type	Type of shaping-rate adjustment performed by the BSR on the scheduler node. The type of adjustment appears as **Adjustment type**, meaning that the configured shaping rate is adjusted by an absolute value as opposed to by a percentage of the configured rate.
Configured shaping rate	The maximum transmission rate on the physical interface as configured by the output traffic-control profile attached to the scheduler node.
Adjustment value	Value of the shaping-rate adjustment information sent by the adjusting application to **cosd**.
Adjustment overhead-accounting mode	Configured shaping mode: **frame** or **cell**.

Sample Output

show class-of-service interface-set

```
user@host> show class-of-service interface-set example-ifset-ge-4/0/0-7
Interface-set: example-ifset-ge-4/0/0-7, Index: 8
Physical interface: ge-4/0/0, Index: 270
Queues supported: 8, Queues in use: 8
  Output traffic control profile: example-tcp-basic-rate, Index: 11395
Adjusting application: ancp LS-0
  Adjustment type: absolute
  Configured shaping rate: 50000000
  Adjustment value: 888000
  Adjustment overhead-accounting mode: cell
```

show class-of-service scheduler-map

Syntax show class-of-service scheduler-map
<*name*>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.

Options **none**—Display all scheduler maps.

name—(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.

Required Privilege Level view

List of Sample Output show class-of-service scheduler-map on page 155

Output Fields Table 15 on page 154 describes the output fields for the **show class-of-service scheduler-map** command. Output fields are listed in the approximate order in which they appear.

Table 15: show class-of-service scheduler-map Output Fields

Field Name	Field Description
Scheduler map	Name of the scheduler map.
Index	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.
Scheduler	Name of the scheduler.
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
Transmit rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword **remainder**, which indicates that the scheduler receives the remaining bandwidth of the interface.
Rate Limit	Rate limiting configuration of the queue. Possible values are **none**, meaning no rate limiting, and **exact**, meaning the queue only transmits at the configured rate.
Maximum buffer delay	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword **remainder** to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	Scheduling priority: **low** or **high**.

Table 15: show class-of-service scheduler-map Output Fields *(continued)*

Field Name	Field Description
Excess priority	Priority of excess bandwidth: **low**, **medium-low**, **medium-high**, **high**, or **none**.
Adjust minimum	Minimum shaping rate for an adjusted queue, in bps.
Adjust percent	Bandwidth adjustment applied to a queue, in percent.
Drop profiles	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
Loss priority	Packet loss priority for drop profile assignment.
Protocol	Transport protocol for drop profile assignment.
Name	Name of the drop profile.

Sample Output

show class-of-service scheduler-map

```
user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

  Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
   Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
   Priority: high
   Drop profiles:
    Loss priority   Protocol    Index    Name
     Low            non-TCP     8724     aa-drop-profile
     Low            TCP         9874     bb-drop-profile
     High           non-TCP     8833     cc-drop-profile
     High           TCP         8484     dd-drop-profile

  Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
   Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
   Priority: high
   Drop profiles:
    Loss priority   Protocol    Index    Name
     Low            non-TCP     8724     aa-drop-profile
     Low            TCP         9874     bb-drop-profile
     High           non-TCP     8833     cc-drop-profile
     High           TCP         8484     dd-drop-profile
```

show class-of-service traffic-control-profile

Syntax show class-of-service traffic-control-profile
<profile-name>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series Routers.

Description For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and Trio MPC/MIC interfaces only, display traffic shaping and scheduling profiles.

(ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.

Options **none**—Display all profiles.

profile-name—(Optional) Display information about a single profile.

Required Privilege Level view

List of Sample Output
show class-of-service traffic-control-profile on page 158
show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC) on page 158
show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces) on page 158

Output Fields Table 16 on page 156 describes the output fields for the **show class-of-service traffic-control-profile** command. Output fields are listed in the approximate order in which they appear.

Table 16: show class-of-service traffic-control-profile Output Fields

Field Name	Field Description
Traffic control profile	Name of the traffic control profile.
Index	Index number of the traffic control profile.
ATM Service	(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: • cbr—Constant bit rate. • rtvbr—Real time variable bit rate. • nrtvbr—Non real time variable bit rate. • ubr—Unspecified bit rate.
Maximum Burst Size	Configured maximum burst size, in cells.
Peak rate	Configured peak rate, in cps.

Table 16: show class-of-service traffic-control-profile Output Fields *(continued)*

Field Name	Field Description
Sustained rate	Configured sustained rate, in cps.
Shaping rate	Configured shaping rate, in bps. NOTE: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
Shaping rate burst	Configured burst size for the shaping rate, in bytes. NOTE: (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.
Shaping rate priority high	Configured shaping rate for high-priority traffic, in bps.
Shaping rate priority medium	Configured shaping rate for medium-priority traffic, in bps.
Shaping rate priority low	Configured shaping rate for low-priority traffic, in bps.
Shaping rate excess high	Configured shaping rate for high-priority excess traffic, in bps.
Shaping rate excess low	Configured shaping rate for low-priority excess traffic, in bps.
Scheduler map	Name of the associated scheduler map.
Delay Buffer rate	Configured delay buffer rate, in bps.
Excess rate	Configured excess rate, in percent or proportion.
Excess rate high	Configured excess rate for high priority traffic, in percent or proportion.
Excess rate low	Configured excess rate for low priority traffic, in percent or proportion.
Guaranteed rate	Configured guaranteed rate, in bps or cps. NOTE: (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values: • **cbr**—Guaranteed rate is equal to the configured peak rate in cps. • **rtvbr**—Guaranteed rate is equal to the configured sustained rate in cps. • **nrtvbr**—Guaranteed rate is equal to the configured sustained rate in cps.
Guaranteed rate burst	Configured burst size for the guaranteed rate, in bytes.
adjust-minimum	Configured minimum shaping rate for an adjusted queue, in bps.

Table 16: show class-of-service traffic-control-profile Output Fields *(continued)*

Field Name	Field Description
overhead accounting mode	Configured shaping mode: **Frame Mode** or **Cell Mode**.
Overhead bytes	Configured byte adjustment value.

Sample Output

show class-of-service traffic-control-profile

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
  Scheduler map: m1
  Delay Buffer rate: 500000
  Guaranteed rate: 1000000

Traffic control profile: Profile2, Index: 57624
  Scheduler map: m2
  Delay Buffer rate: 600000
  Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
  Scheduler map: m3
  Delay Buffer rate: 800000
  Guaranteed rate: 3000000
 .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
  Scheduler map: m4
  Delay Buffer rate: 750000
  Guaranteed rate: 4000000
..adjust-minimum 20000000
```

show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
  ATM Service: RTVBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
  Shaping rate: 1000 cps
  Shaping rate burst: 500 cells
  Delay Buffer rate: 2000 cps
  Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
  ATM Service: UBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
```

show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```
user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
  ATM Service: RTVBR
  Shaping rate: 2000 cps
```

```
Shaping rate burst: 200 cells
Scheduler map: <default>
Delay Buffer rate: 1000 cps
Guaranteed rate: 1700 cps
```

show class-of-service adjustment-control-profile

Syntax show class-of-service adjustment-control-profile <*profile-name*>

Release Information Command introduced in Junos OS Release 13.1 for MX Series Routers.

Description For MPC/MIC interfaces only, display the adjustment control profiles.

Options **none**—Display all profiles.

profile-name—(Optional) Display information about a single profile.

Required Privilege Level view

Related Documentation
- Verifying the CoS Adjustment Control Profile Configuration on page 119

List of Sample Output show class-of-service adjustment-control-profile on page 161

Output Fields Table 17 on page 160 describes the output fields for the **show class-of-service adjustment-control-profile** command. Output fields are listed in the approximate order in which they appear.

Table 17: show class-of-service adjustment-control-profile Output Fields

Field Name	Field Description
Name	Name of the adjusting application. Possible values: • RADIUS-CoA—RADIUS CoA aapplication. • ANCP—ANCP application. • PPPoE IA tags—PPPoE IA tag application.
Priority	Priority of the adjusting application. Possible values are1 through 10; 1 being the highest priority. The lower the priority value, the higher the priority
Algorithm	Algorithm the adjusting application uses to make adjustments. • adjust-never—Never perform rate adjustments. • adjust-always—Adjust the shaping rate unconditionally. • adjust-less—Adjust the shaping rate if it is less than the configured value. • adjust-less-or equal—Adjust the shaping rate if it is less than or equal to the configured value. • adjust-greater—Adjust the shaping rate if it is greater than the configured value. • adjust-greater-or-equal—Adjust the shaping rate if it is greater than or equal to the configured value.

Sample Output

show class-of-service adjustment-control-profile

```
user@host> show class-of-service adjustment-control-profile

  name: ANCP, priority: 1, algorithm: less
  name: RADIUS CoA, priority: 1, algorithm: always
  name: PPPoE IA tags, priority: 2, algorithm: less
```

PART 4

Troubleshooting

CHAPTER 13

Acquiring Troubleshooting Information

This chapter describes messages with the **COSD** prefix. They are generated by the class-of-service (CoS) process (cosd), which enables the routing platform to provide different levels of service to applications based on packet classifications.

COSD_AGGR_CONFIG_INVALID

System Log Message Error: Cannot have config *error-message interface-name*

Description The class-of-service (CoS) process (cosd) did not apply the config on this interface because it was not valid in this case.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause One possible cause is if any Class-of-Service is configured on an interface which is a part of an aggregated interface

Action Remove or change the config from/on the interface.

COSD_CHASSIS_SCHED_MAP_INVALID

System Log Message Chassis scheduler map incorrectly applied to interface *interface-name*: *error-message*

Description The class-of-service (CoS) process (cosd) did not apply a chassis scheduler map to the indicated interface, because the configuration used to apply the scheduler map was invalid.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause One possible cause is that the chassis scheduler map is applied to a specific interface. For most interface types, a scheduler map must be applied to all interfaces on the PIC;

therefore, a wildcard must be used to specify the interfaces. One exception to this rule is the Gigabit Ethernet IQ PIC.

Action Correct the configuration used to apply the chassis scheduler map to the interface.

COSD_CLASSIFIER_NO_SUPPORT_LSI

System Log Message Cannot support classifier type *classifier-type* on lsi interface *interface-name*

Description The Differentiated Services code point (DSCP) classifier and the 802.1p classifier are only supported on I-Chip based Flexible PIC Concentrators (FPCs).

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the DSCP or the 802.1p classifier configuration from the routing instance

COSD_CLASS_8021P_UNSUPPORTED

System Log Message ieee-802.1 classifier is not valid on interface *interface-name*

Description The IEEE 802.1p classifier is not supported on the indicated interface.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Remove the 802.1p classifier configuration from the interface, or configure an interface encapsulation type that supports 802.1p classifiers.

COSD_CLASS_NO_SUPPORT_IFD

System Log Message BA/Fixed Classifier or Rewrite on Physical Interface is not allowed when ethernet switching family is configured: interace *interface-name*

Description The Rewrite is not supported on this interface when ethernet switching is enabled

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the classifier configuration from the interface, instead apply it on the logical interface where ethernet switching family is enabled

Action Remove the Rewrite configuration from the interface, instead apply it on the logical interface where ethernet switching family is enabled

COSD_CLASS_NO_SUPPORT_L3_IFL

System Log Message BA/Fixed Classifier or Rewrite config is not allowed on logical interface (*interface-name*) with inet/inet6 family

Description The Rewrite is not supported on this logical interface

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the classifier configuration from the logical interface, instead apply it on the main interface if inet/inet6 is configured on one of its logical interfaces

Action Remove the Rewrite configuration from the logical interface, instead apply it on the main interface if inet/inet6 is configured on one of its logical interfaces

COSD_CONF_OPEN_FAILURE

System Log Message Unable to open: *filename*, using default CoS forwarding classes, do 'commit full' in cli to avoid this message

Description The class-of-service (CoS) process (cosd) could not read configuration data.

Type Error: An error occurred

Severity error

Facility ANY

Cause All of the following resons: mgd -I fails after upgrade-the file cosd.conf does not exist and is not created because of the mgd -I failure The first commit is 'commit' and not 'commit full'-the file cosd.conf does not commit and is not created automatically [class-of-service forwarding-classes] does not exist-the file cosd.conf does not get exported with plain 'commit'

Action Do a 'commit full'

COSD_DB_OPEN_FAILED

System Log Message Unable to open configuration database: *error-message*(*name*)

Description The class-of-service (CoS) process (cosd) could not read configuration data for the indicated reason.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause The specified database does not exist

Action Contact your technical support representative.

COSD_EXACT_RATE_UNSUPP_INTERFACE

System Log Message Unable to apply scheduler map *scheduler-map* to interface *interface-name* because it does not support exact-rate transmission

Description The class-of-service (CoS) process (cosd) did not apply the indicated scheduler map to the indicated interface, because a scheduler named in the scheduler map specifies exact transmission rate. The interface is housed on a type of PIC that does not support exact transmission rate, such as an IQ2 PIC. In terms of configuration, the 'exact' statement is included in the scheduler definition at the [edit class-of-service schedulers <scheduler-name> transmit-rate (<rate> | percent <percentage>)] hierarchy level. The scheduler is included in the scheduler map that is applied to the interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the 'exact' statement from the scheduler in the scheduler map applied to the interface.

COSD_EXACT_RATE_UNSUPP_SESSION

System Log Message Unable to apply CoS to L2TP session *session-id*, because scheduler map *scheduler-map* specifies exact rate transmission

Description The class-of-service (CoS) process (cosd) did not apply CoS settings to the indicated Layer 2 Tunneling Protocol (L2TP) session, because the scheduler map specified by the RADIUS server for the session is configured for exact transmission rate. Exact transmission rate is not supported for L2TP sessions on the type of PIC that houses the interface, such as an IQ2 PIC. In terms of configuration, the 'exact' statement is included in a scheduler definition at the [edit class-of-service schedulers <scheduler-name> transmit-rate (<rate> | percent <percentage>)] hierarchy level. The scheduler is included in a scheduler map that is associated with a traffic control profile. The traffic control profile is named by an attribute in the RADIUS server's configuration file, which makes the profile apply to the session.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the 'exact' statement from the scheduler in the scheduler map applied to the session.

COSD_EXP_RW_L2_IFL_NOT_SUPPORTED

System Log Message EXP Rewrite on IFL is not allowed when ethernet switching family is configured: interace *interface-name*

Description EXP rewrite is not supported on this logical interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the exp rewrite configuration from the logical interface.

COSD_FRAGMENTATION_MAP_CONFLICT

System Log Message Interface *compression-device* matches wildcard *wildcard-interface-name*, but fragmentation map *fragmentation-map* was not applied because interface is compression device for link interface *link-interface-name*

Description The indicated fragmentation map is normally applied to interfaces that match the indicated wildcard. The class-of-service (CoS) process (cosd) did not apply the fragmentation map to the indicated interface, even though it matches the wildcard, because the interface is acting as a compression device for the indicated link interface.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Correct the configuration of the fragmentation map.

COSD_HIGH_PRIO_QUEUES_INTERFACE

System Log Message Unable to apply scheduler map *scheduler-map* to interface *interface-name*, because multiple schedulers in map have "high,""medium-high," or "strict-high" priority

Description The class-of-service (CoS) process (cosd) did not apply the indicated scheduler map to the indicated interface, because the map includes more than one scheduler that has high, medium-high, or strict-high priority. For interfaces that are housed by certain PICs, such as an IQ2 PIC, the scheduler map can include only one scheduler that specifies one of those three priority levels. In terms of configuration, the 'priority' statement at the [edit class-of-service schedulers <scheduler-name>] hierarchy level has the value 'high, ' 'medium-high, ' or 'strict-high' for more than one of the schedulers in the map.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Correct the configuration so that the scheduler map includes only one scheduler with high, medium-high, or strict-high priority.

COSD_HIGH_PRIO_QUEUES_SESSION

System Log Message Unable to apply CoS to L2TP session *session-id*, because multiple schedulers in scheduler map *scheduler-map* have "high,""medium-high," or "strict-high" priority

Description The class-of-service (CoS) process (cosd) did not apply CoS settings to the indicated Layer 2 Tunneling Protocol (L2TP) session because the scheduler map specified by the RADIUS server for the session includes more than one scheduler that has high, medium-high, or strict-high priority. For interfaces that are housed by certain Physical Interface Cards (PICs), such as an IQ2 PIC, the scheduler map can include only one scheduler that specifies one of those three priority levels. In terms of configuration, the 'priority' statement at the [edit class-of-service schedulers <scheduler-name>] hierarchy level has the value 'high, ' 'medium-high, ' or 'strict-high' for more than one of the schedulers in the map. The map is associated with a traffic control profile that is named by an attribute in the RADIUS server's configuration file, which makes the profile apply to the session.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Correct the configuration so that the scheduler map includes only one scheduler with high, medium-high, or strict-high priority.

COSD_IFD_OUTPUT_SHAPING_RATE_ERR

System Log Message Traffic shaping not supported on interface device *interface-name*

Description The class-of-service (CoS) process (cosd) did not apply the shaping rate that is configured for the indicated interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause Shaping rate is valid only for interfaces housed by IQ and IQ2 PICs, and the interface is on a different type of PIC.

Action Remove the shaping rate configuration from the interface.

COSD_IFD_SHAPER_ERR

System Log Message port shaper not allowed on interface *interface-name*

Description The non-queuing dense port concentrators (DPCs) did not support the specified shaping rate.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause The port shaper was not supported on the non-queuing DPCs.

Action Remove the shaping rate configuration from the interface.

COSD_INTERFACE_NO_MEDIA

System Log Message Unable to obtain media information for interface *interface-name*

Description The message sent by the kernel for the indicated interface did not include required media information.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause An internal software failure occurred.

Action Contact your technical support representative.

COSD_L2TP_COS_NOT_CONFIGURED

System Log Message Unable to apply CoS to L2TP session *session-id* because session-aware CoS is not enabled for interface *interface-name*

Description The class-of-service (CoS) process (cosd) did not apply CoS settings to the indicated Layer 2 Tunneling Protocol (L2TP) session on the indicated interface, because the interface is not configured to support session-aware CoS for L2TP. In terms of configuration, the 'per-session-scheduler' statement is not included at the [edit interfaces <interface-name> unit <logical-unit-number>] hierarchy level.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Include the 'per-session-scheduler' statement in the configuration for the interface.

COSD_L2TP_COS_NOT_SUPPORTED

System Log Message Unable to apply CoS to L2TP session *session-id* on interface *interface-name*: it does not support CoS

Description The class-of-service (CoS) process (cosd) did not apply CoS settings to the indicated Layer 2 Tunneling Protocol (L2TP) session on the indicated interface. The interface is configured to support session-aware CoS for L2TP, but is not on a PIC that supports that feature, such as an IQ2 PIC. In terms of configuration, the 'per-session-scheduler'

statement is included at the [edit interfaces <interface-name> unit <logical-unit-number>] hierarchy level.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Determine whether the interface is on an PIC that supports session-aware CoS; if not, remove the 'per-session-scheduler' statement.

COSD_L2TP_SHAPING_NOT_CONFIGURED

System Log Message Unable to apply CoS to L2TP session *session-id* because session-aware shaping is not enabled for interface *interface-name*

Description The class-of-service (CoS) process (cosd) did not apply CoS settings to the indicated Layer 2 Tunneling Protocol (L2TP) session on the indicated interface, because session-aware traffic shaping for L2TP is not configured on the PIC that houses the interface. In terms of configuration, the 'session-shaping' statement is not included at the [edit chassis fpc <slot-number> pic <pic-number> traffic-manager mode] hierarchy level.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Include the 'session-shaping' statement in the configuration for the PIC.

COSD_LARGE_DELAY_BUFFER_INVALID

System Log Message Error for interface *interface-name error-message*

Description The class-of-service (CoS) process (cosd) did not apply the large delay buffer setting that is configured for the indicated interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause The interface is not housed on one of the PIC types that support large delay buffer.

Action Remove the large delay buffer configuration from the interface.

COSD_MALLOC_FAILED

System Log Message malloc failed: *error-message*

Description The class-of-service (CoS) process (cosd) could not dynamically allocate memory for the indicated reason.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause A software bug caused a memory leak or the Routing Engine did not have sufficient memory.

Action Contact your technical support representative. For more information, see http://kb.juniper.net/InfoCenter/index?page=content&id=KB18862.

COSD_MAX_FORWARDING_CLASSES_ABC

System Log Message exceeding max 4 forwarding-class support.

Description User configuration exceeds the maximum number of forwarding class that is supported.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Configure only four forwarding classes

COSD_MPLS_DSCP_CLASS_NO_SUPPORT

System Log Message Cannot support MPLS DSCP classifier on ifl *interface-name*

Description The MPLS Differentiated Services code point (DSCP) classifier is only supported on I-Chip based Flexible PIC Concentrators (FPCs). It is not supported on Q2 PICs.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the MPLS DSCP classifier configuration from the logical interface.

COSD_MULTILINK_CLASS_CONFLICT

System Log Message Fragmentation map *fragmentation-map* for wildcard *wildcard-interface-name* specified multilink class *class-name* for queue *queue-number* on interface *interface-name*, which exceeds configured limit of *limit*

Description The indicated fragmentation map is normally applied to interfaces that match the indicated wildcard, and specifies the indicated multilink class setting for queues on those interfaces. The class-of-service (CoS) process (cosd) did not apply the fragmentation

map to the indicated interface, even though it matches the wildcard, because the setting in the map exceeds the indicated class limit, which is configured on the interface itself.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Correct the configuration so that the multilink class setting in the fragmentation map does not exceed the class limit for the interface.

COSD_NULL_INPUT_ARGUMENT

System Log Message NULL input argument : *error-message*

Description The pointer that was passed to this function was NULL.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Contact your technical support representative.

COSD_OUT_OF_DEDICATED_QUEUES

System Log Message Queue usage count for interface *interface-name* is at *percentage-value* percent

Description The class-of-service (CoS) process (cosd) is running out of dedicated queues.

Type Event: This message reports an event, not an error

Severity warning

Facility LOG_DAEMON

COSD_RATE_LIMIT_INVALID

System Log Message Unable to apply scheduler map *scheduler-map* to interface *interface-name*. *description*.

Description The class-of-service (CoS) process (cosd) did not apply the indicated scheduler map to the indicated interface, because the number of rate limited queues in the scheduler map exceeded the limit supported by this interface or the priority is not supported. The interface is housed in a type of PIC that does not support the number of configured rate limited queues or the priority is not supported. In terms of configuration, the 'rate-limit' statement is included in the scheduler definition at the [edit class-of-service schedulers <scheduler-name> transmit-rate <rate> | percent <percentage>] hierarchy level. The scheduler is included in the scheduler map applied to the interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Either limit the number of rate-limited schedulers in this scheduler map to the allowed maximum for this PIC and interface type or check the allowed priority for rate-limited queues

COSD_RATE_LIMIT_NOT_SUPPORTED

System Log Message Unable to apply scheduler map *scheduler-map* to interface *interface-name* because it does not support rate limiting

Description The class-of-service (CoS) process (cosd) did not apply the indicated scheduler map to the indicated interface, because a scheduler named in the scheduler map is configured for rate limiting. The interface is housed in a type of PIC that does not support rate limiting. In terms of configuration, the 'rate-limit' statement is included in the scheduler definition at the [edit class-of-service schedulers <scheduler-name> transmit-rate <rate> | percent <percentage>] hierarchy level. The scheduler is included in the scheduler map applied to the interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove the 'rate-limit' statement from the scheduler in the scheduler map applied to the interface.

COSD_REWRITE_RULE_LIMIT_EXCEEDED

System Log Message Number of rewrite rules applied to interface *interface-name* exceeds limit (*maximum-value*)

Description The class-of-service (CoS) process (cosd) determined that the number of rewrite rules applied to the indicated interface exceeds the indicated limit for the interface. In terms of configuration, too many rewrite rules are included at the [edit class-of-service interfaces <interface-name> unit <logical-unit-number> rewrite-rules] hierarchy level.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Remove rewrite rules from the configuration for the interface.

COSD_RL_IFL_NEEDS_SHAPING

System Log Message "rate-limit" configured in *scheduler-map*, but ifl *interface-name* does not have output shaper configured. It will use the ifd-shaping rate/ifd-transmit rate for implementation of rate-limit.

Description The 'rate-limit' statement is configured in one or more schedulers that are part of the indicated scheduler map. In order to apply this scheduler map to the indicated interface,

output shaping rate should be configured on the interface. Since no output shaping rate is configured, the transmit rate or shaping rate of the parent interface will be used instead.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Configure output shaping rate for the indicataed interface

COSD_SCHEDULER_MAP_CONFLICT

System Log Message Forwarding classes "*first-forwarding-class*" and "*second-forwarding-class*" in scheduler map *scheduler-map* both map to queue *queue-number*

Description Both of the indicated forwarding classes, which are defined in the indicated scheduler map, map to the same indicated queue. The double mapping is invalid.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Map only one forwarding class to the queue.

COSD_SCHED_AVG_CONST_UNSUPPORTED

System Log Message Averaging constant not supported on interface *interface-name*. Value set in scheduler-map *scheduler-map* (scheduler *name*) will be ignored.

Description Configuring averaging constant is not supported on the indicated interface. Value set in the indicated scheduler will be ignored.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Remove the averaging-constant configuration from the indicated scheduler.

COSD_SCHED_MAP_GROUP_CONFLICT

System Log Message Interface *interface-name* cannot be bound to scheduler-map *scheduler-map*. It will be bound to default scheduler-map

Description Interfaces belonging to a group cannot be bound to different scheduler maps. They will be bound to the default scheduler map.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Map only one scheduler map to all the interfaces of a group.

COSD_SHAPER_GROUP_CONFLICT

System Log Message Interface *interface-name* cannot be bound to configured shaping-rate. It will be bound to default rate

Description Interfaces belonging to a group cannot be bound to different shaping rates. They will be bound to the default shaping rate.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Map only one shaping rate to all interfaces of a group.

COSD_STREAM_IFD_CREATE_FAILURE

System Log Message Unable to create special master interface device for *interface-name*

Description The class-of-service (CoS) process (cosd) could not create the indicated internal interface device, which it needs for application of a chassis scheduler map.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause An internal software failure occurred.

Action Contact your technical support representative.

COSD_TIMER_ERROR

System Log Message Unable to set retry timer for rtsock write operation: *error-message*

Description The class-of-service (CoS) process (cosd) used a routine from the rtsock library to write to the kernel, but the kernel did not accept the request. The cosd process could not set the retry timer for the request, for the indicated reason.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Cause An internal software failure occurred.

Action Contact your technical support representative.

COSD_TRICOLOR_ALWAYS_ON

System Log Message tri-color is always enabled in this platform. There is no need to explicitly set it.

Description Tri-color marking is always enabled on this platform. There is no need to explicitly set it.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Action Remove the 'tri-color' configuration statement

COSD_TRICOLOR_NOT_SUPPORTED

System Log Message Unable to apply scheduler *scheduler-map* to interface *interface-name*, because it does not support tricolor marking

Description The class-of-service (CoS) process (cosd) did not apply the indicated scheduler map to the indicated interface, because a scheduler included in the map specifies a packet loss priority (PLP) that is supported only with tricolor marking (TCM). The interface does not support TCM, either because TCM is not enabled or the interface is on a router that does not support TCM. In terms of configuration, the value 'medium-high' or 'medium-low' is specified for the 'loss-priority' statement in a scheduler definition at the [edit class-of-service schedulers <scheduler-name> drop-profile-map] hierarchy level. The scheduler is included in the scheduler map applied to the interface, but the 'tri-color' statement is either not included at the [edit class-of-service] hierarchy level, or is not supported.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Change the value of the 'loss-priority' statement in the scheduler or include the 'tri-color' statement to enable TCM on the router.

COSD_TX_QUEUE_RATES_TOO_HIGH

System Log Message Unable to apply scheduler map *scheduler-map* to interface *interface-name*: sum of scheduler transmission rates exceeds interface shaping or transmission rate

Description The class-of-service (CoS) process (cosd) did not apply the indicated scheduler map to the indicated interface, because the sum of the queue transmission rates defined in the schedulers in the scheduler map exceeds the shaping or transmission rate for the interface. In terms of configuration, the 'transmit-rate' statement is specified for each scheduler at the [edit class-of-service schedulers <scheduler-name>] hierarchy level. The sum of the configured transmission rates exceeds the transmission or shaping rate of the interface.

Type Error: An error occurred

Severity error

Facility LOG_DAEMON

Action Decrease the value of one or more 'transmit-rate' statements so that the sum is less than the interface transmission or shaping rate.

COSD_UNKNOWN_CLASSIFIER

System Log Message classifier type *classifier-type* is invalid

Description The class-of-service (CoS) process (cosd) did not recognize the indicated classifier type from the rtsock library.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Cause An internal software failure occurred.

Action Contact your technical support representative.

COSD_UNKNOWN_REWRITE

System Log Message rtsock rewrite type *type* is invalid

Description The class-of-service (CoS) process (cosd) did not recognize the indicated rewrite type from the rtsock library.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Cause An internal software failure occurred.

Action Contact your technical support representative.

COSD_UNKNOWN_TRANSLATION_TABLE

System Log Message rtsock translation table type *translation-table-type* is invalid

Description The class-of-service (CoS) process (cosd) did not recognize the indicated translation table type from the rtsock library.

Type Error: An error occurred

Severity warning

Facility LOG_DAEMON

Cause An internal software failure occurred.

Action For more information, see http://kb.juniper.net/InfoCenter/index?page=content&id=KB18866.

Collecting Subscriber Access Logs Before Contacting Juniper Technical Support

Problem When you experience a subscriber access problem in your network, we recommend that you collect certain logs before you contact Juniper Technical Support. This topic shows you the most useful logs for a variety of network implementations. In addition to the relevant log information, you must also collect standard troubleshooting information and send it to Juniper Technical Support in your request for assistance.

Solution To collect standard troubleshooting information:

- Redirect the command output to a file.

  user@host> **request support information | save rsi-1**

To configure logging to assist Juniper Technical Support:

1. Review the following blocks of statements to determine which apply to your configuration.

   ```
   [edit]
   set system syslog archive size 100m files 25
   set system auto-configuration traceoptions file filename
   set system auto-configuration traceoptions file filename size 100m files 25
   set protocols ppp-service traceoptions file filename size 100m files 25
   set protocols ppp-service traceoptions level all
   set protocols ppp-service traceoptions flag all
   set protocols ppp traceoptions file filename size 100m files 25
   set protocols ppp traceoptions level all
   set protocols ppp traceoptions flag all
   set protocols ppp monitor-session all
   set interfaces pp0 traceoptions flag all
   set demux traceoptions file filename size 100m files 25
   set demux traceoptions level all
   set demux traceoptions flag all
   set system processes dhcp-service traceoptions file filename
   set system processes dhcp-service traceoptions file size 100m
   set system processes dhcp-service traceoptions file files 25
   set system processes dhcp-service traceoptions flag all
   set class-of-service traceoptions file filename
   set class-of-service traceoptions file size 100m
   set class-of-service traceoptions flag all
   set class-of-service traceoptions file files 25
   set routing-options traceoptions file filename
   set routing-options traceoptions file size 100m
   set routing-options traceoptions flag all
   set routing-options traceoptions file files 25
   set interfaces traceoptions file filename
   set interfaces traceoptions file size 100m
   set interfaces traceoptions flag all
   set interfaces traceoptions file files 25
   set system processes general-authentication-service traceoptions file filename
   set system processes general-authentication-service traceoptions file size 100m
   set system processes general-authentication-service traceoptions flag all
   set system processes general-authentication-service traceoptions file files 25
   ```

2. Copy the relevant statements into a text file and modify the log filenames as you want.
3. Copy the statements from the text file and paste them into the CLI on your router to configure logging.
4. Commit the logging configuration to begin collecting information.

NOTE: The maximum file size for DHCP local server and DHCP relay log files is 1 GB. The maximum number of log files for DHCP local server and DHCP relay is 1000.

BEST PRACTICE: Enable these logs only to collect information when troubleshooting specific problems. Enabling these logs during normal operations can result in reduced system performance.

Related Documentation

- *Compressing Troubleshooting Logs from /var/logs to Send to Juniper Technical Support*

PART 5

Index

Index

Symbols

A

B

C

U

V