Performance Management — Quality of Service
Avaya Secure Router 2330/4134

Release 10.3.5
NN47263-601
Issue 04.02
August 2013

© 2013 Avaya Inc.

All Rights Reserved.

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya generally makes available to users of its products. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on its hardware and Software (“Product(s)”). Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this Product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya. “Software” means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products or pre-installed on hardware products, and any upgrades, updates, bug fixes, or modified versions.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or hardware provided by Avaya. All content on this site, the documentation and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software that may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). Information regarding distributed Linux OS source code (for those Products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the Documentation or on Avaya’s website at: http://support.avaya.com/Copyright. You agree to the Third Party Terms for any such Third Party Components.

Note to Service Provider

The Product may use Third Party Components that have Third Party Terms that do not allow hosting and may need to be independently licensed for such purpose.

Preventing Toll Fraud

“Toll Fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation and Product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation and Product(s) should be construed as granting, by implication, estoppel, or otherwise,

any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product notices and articles, or to report a problem with your Avaya product. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com, scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Chapter 1: Introduction

Purpose

This document describes the operation and configuration of the QoS features on the Avaya Secure Router 2330/4134 (SR2330/4134).

Related resources

Documentation

See the *Avaya Secure Router 2330/4134 Documentation Roadmap*, NN47263-103, for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the Web site at http://avaya-learning.com.

Avaya Mentor videos

Avaya Mentor is an Avaya-run channel on YouTube that includes technical content on how to install, configure, and troubleshoot Avaya products.

Go to http://www.youtube.com/AvayaMentor and perform one of the following actions:

- Enter a key word or key words in the Search Channel to search for a specific product or topic.
- Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the site.

Support

Visit the Avaya Support website at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Chapter 2: New in this release

There is no new content added to *Performance Management — Quality of Service Avaya Secure Router 2330/4134* (NN47263-601) for Release 10.3.5.

Chapter 3: QoS fundamentals

QoS on the Avaya Secure Router 2330/4134 is logically separated into three areas:

- Chassis QoS
- Ethernet Module QoS
- SLA

The following figure shows the logical view of the division of operations between the two main QoS components: Chassis QoS and Ethernet Module QoS.

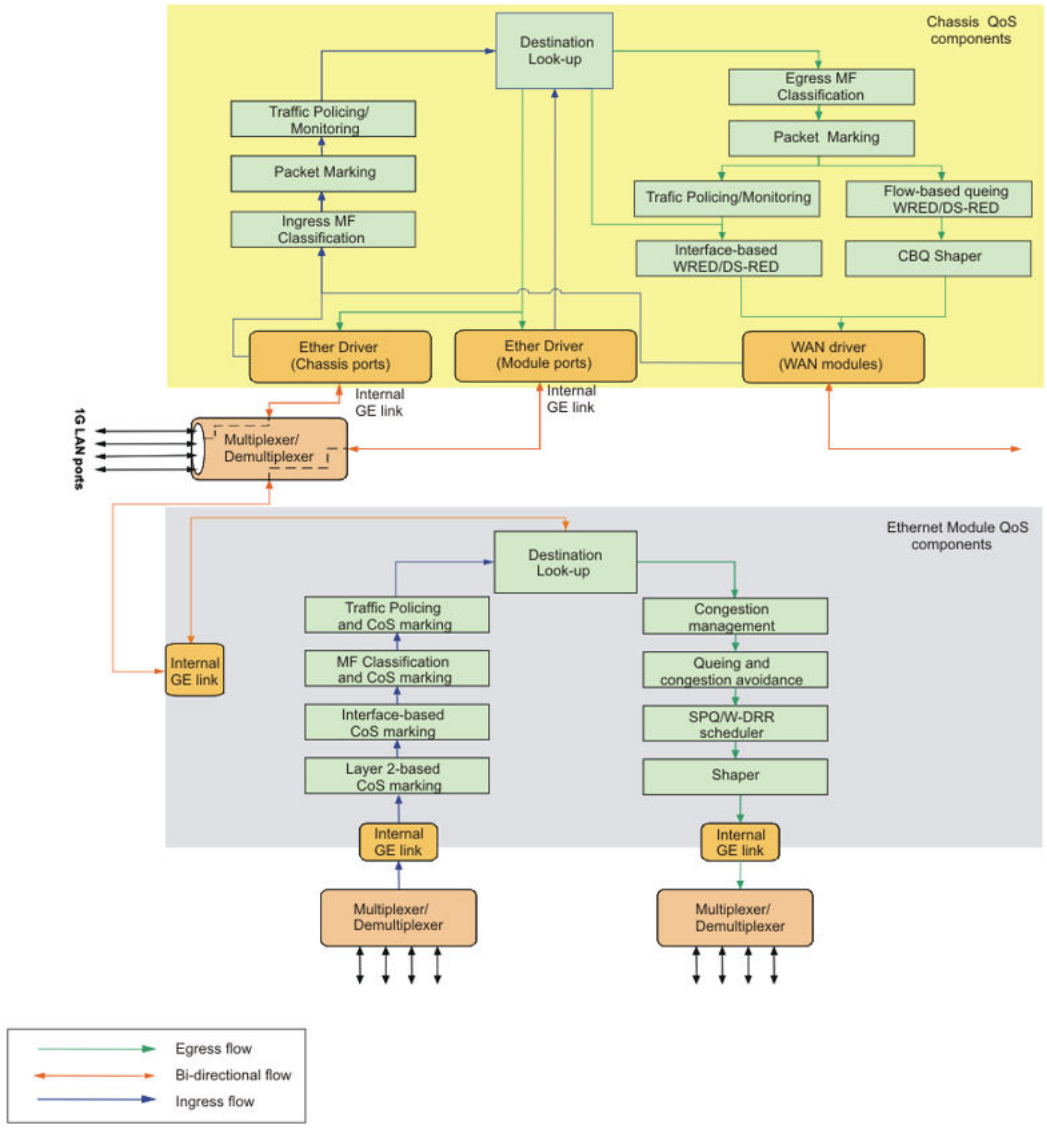

Figure 1: Complete view of QoS processing

Chassis QoS

Chassis QoS provides the QoS for all ingress or egress traffic on all of the on-board chassis Ethernet ports (0/1-0/4), as well as all WAN interfaces.

The Chassis QoS is also applied for logical chassis interfaces such as tunnels interfaces, Crypto interfaces (for IPSec VPN), and VLAN interfaces.

It can also provide egress QoS for traffic from the ethernet modules that is destined to exit one of the chassis interfaces.

On the SR2330, Chassis QoS refers to QoS for the following:

- All non-Ethernet interfaces.
- Egress traffic on Ethernet interfaces.
- WAN interfaces and logical interfaces including tunnels, crypto interfaces, and VLAN interfaces.

Ethernet Module QoS

All Ethernet modules support Layer 2 switching and Layer 3 routing functionality. In addition, the Ethernet modules provide an on-board QoS functionality that is separate from the Chassis QoS system. Provided that packets entering the Ethernet modules are not destined to a chassis interface (WAN interface or chassis Ethernet ports), then QoS processing is performed solely on the Ethernet modules. Packets that travel in either direction between the Ethernet module ports and the chassis ports are subject to QoS processing from both QoS systems.

On the SR2330, Ethernet module QoS refers to the QoS on all front-panel Ethernet ports (FE and GE).

SLA

SLA is an additional QoS component that serves as a performance monitoring and measuring tool that you can use to monitor and measure network service performance between two nodes.

Chapter 4: Chassis QoS fundamentals

The following figure shows the overview of features available with Chassis QoS flows.

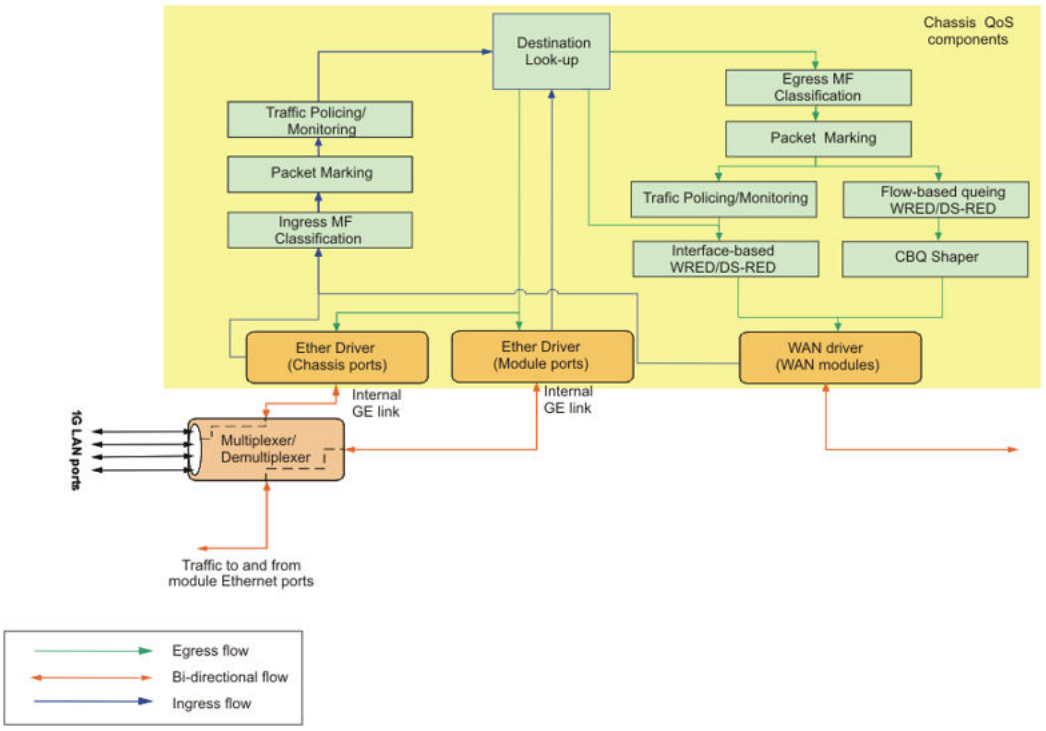

Figure 2: Overview of Chassis QoS operation

The following sections provide details for each of the Chassis QoS features.

Multifield traffic classification

Multifield traffic classification is the first step in the Chassis QoS process. With multifield traffic classification, you can identify and segregate traffic that matches particular criteria to classify aggregate traffic into separate flows or traffic classes.

The traffic classification can be configured for inbound and outbound traffic on any chassis interface. However, the available QoS options differ on inbound and outbound flows.

Multifield traffic classification uses a hierarchical classification tree approach. In this approach, each packet is passed through multiple levels of classification. At each level, the packet is classified based on only one classification attribute (header field). Each level classifies the packet based on a different classification attribute. Essentially, the multifield classification process is broken down into simpler classification steps, each involving a lookup based on a single field at each level of the hierarchical tree. The packet is considered to be successfully classified into a flow when the classification process ends on a leaf node of the hierarchical classification tree.

Once traffic is classified into a leaf node, you can then configure desired QoS parameters to define the forwarding treatment for the packets matching that node. The available QoS options include policing, scheduling, monitoring, and so on.

Policy map

A policy map consists of a hierarchy of traffic classes or class maps that classify and further sub-classify packets into finer flows as shown in the figure below.

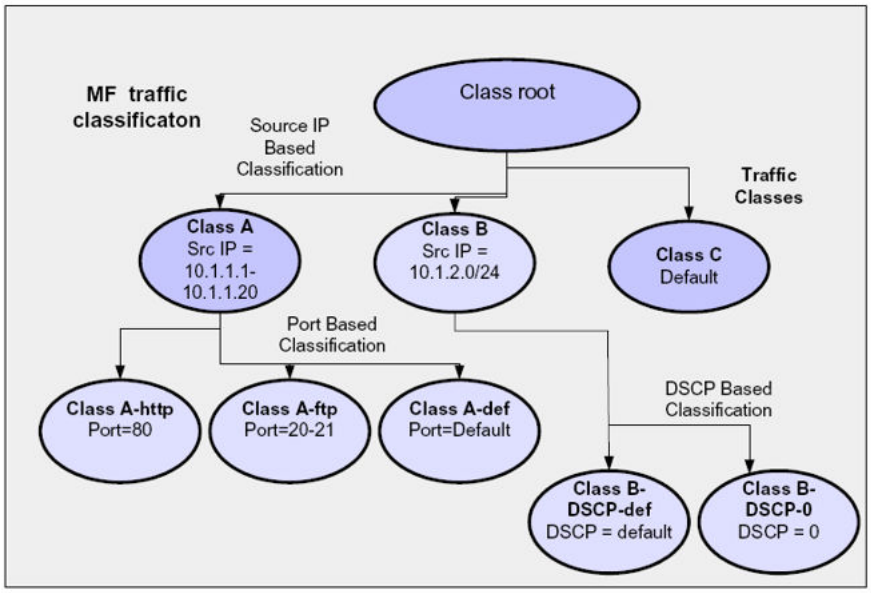

Figure 3: Policy map

The root of a policy map is the interface to which it is mapped. All traffic on the interface is matched to the root class.

Classification starts at the root class and ends at the leaf class. The leaf class defines the traffic class for which QoS parameters can be configured, whereas the intermediate, non-leaf classes aid primarily in the classification process.

For example, the hierarchical classification shown in the preceding figure defines the following traffic flows:

Class name	Classification attribute	Classification attribute value
Class A	Source IP	10.1.1.1-10.1.1.20
Class A-http	Source IP and port	10.1.1.1-10.1.1.20 and Port = 80
Class A-ftp	Source IP and port	10.1.1.1-10.1.1.20 and Ports = 20-21
Class A-def	Source IP and port	10.1.1.1-10.1.1.20 AND Ports = default
Class B	Source IP	10.1.2.0/24
Class B-def	Source IP and DSCP	10.1.2.0/24 and DSCP value default
Class B-0	Source IP and DSCP	10.1.2.0/24 and DSCP value 0
Class C Interface-default	Default	Default

To create a new node or class within a policy map, you must specify the class name of the new class, the class name of the parent class, and a classification match rule.

To add classes at the first level, you must specify the name of the root class (root) as the parent class.

Default class

With certain interfaces, a default class is automatically created for inbound or outbound traffic on the interface. If a packet does not match any of your user-defined classes, then it is classified to the interface default class (Class C in the preceding figure).

You can also create additional default classes at each level of the hierarchy (for example, Class A-def, Class-B-def in the previous table). Any traffic that does not get classified to a class at a particular level of hierarchy is routed to the default class at that level. If you have not created a default class for that level, then the traffic is routed to the interface default class.

The following table shows which interfaces are automatically configured with default classes.

Table 1: Interfaces with default classes

Interface	Inbound default class	Outbound default class
Ethernet	Yes	Yes
HDLC	No	Yes
PPP	No	Yes
Frame Relay	No	Yes
VLAN	No	No
Tunnel	No	No

Interface	Inbound default class	Outbound default class
Crypto	No	No

For Ethernet, PPP, and HDLC interfaces, the default class is configured with CBQ using a Committed Rate (CR) of 1% and a Peak Rate (PR) of 100%.

For Frame Relay interfaces, an outbound default class is automatically created for each PVC. The default class is configured with a CR that is initially set to 100% of the PVC CIR (shaping parameter) and a PR that is initially set to 100% of the bundle bandwidth. When you add a class to the PVC, the configured CR of the new class is borrowed from the outbound default class of that particular PVC. While the CR of the outbound default class is reduced accordingly to meet the requirements of each new class, the default class always maintains a minimum of 1% of the PVC CIR.

Classification attributes

The field to be matched in a particular match rule is referred to as the classification attribute. For example, in the preceding figure, for Class A, Source IP is the classification attribute.

At any parent node in the class map, all children must have the same classification attribute. No heterogeneous classification attributes are allowed. To specify a different classification attribute, you must create additional child classes at the next level.

Under any one parent node, when you define the match rule for the first child class, any subsequent child classes under the same parent must use the same classification attribute for a match rule.

For example, in the preceding figure, if you define class A as a function of the source IP, then all other children under the root parent class must also be defined as a function of the source IP.

The supported classification attributes are:

- IPv4 or IPv6 source address (or prefix)
- IPv4 or IPv6 destination address (or prefix)
- Source/destination TCP/UDP ports
- IPv4 or IPv6 DiffServ code point (DSCP)
- IPv4 Type Of Service (TOS)
- IPv4 precedence
- IPv4 protocol type
- IPv6 Traffic Class
- IPv6 next header (same as protocol type of IPv4)

- IPv6 flow label
- MPLS label
- MPLS EXP bits
- VLAN ID
- IEEE 802.1p User Priority values (Ethernet CoS attribute)

Important:

With Q-in-Q stacked VLAN packets, classification using VLAN ID or user priority is based only on the outer VLAN tag.

Classification value ranges

For each of the above classification attributes, you can define multiple values to define more aggregate or coarser flows. You can specify classification values either as a range of values or as multiple non-contiguous values. There is no restriction on the number of classification values. For example, any number of IP addresses can be assigned to a traffic class for creating a match rule.

PHB group code points

To avoid packet re-ordering for a given traffic flow from source to destination, you cannot spread the AF1x PHB group code points across nodes. The same restriction applies for AF2x, AF3x, and AF4x group code points.

Policy map and QoS actions

After you have a leaf class defined, you can associate QoS parameters to that leaf class.

The possible QoS features that you can apply to the leaf class are as follows:

Inbound

- CoS marking
- policy-based redirect
- policing (srTCM/trTCM)
- monitoring

Important:

On inbound traffic, policing and monitoring are mutually-exclusive features. When you enable one features, the other feature is automatically disabled.

Outbound

- CoS marking
- CBQ (queuing, scheduling and shaping)
- policing (srTCM/trTCM)
- monitoring
- congestion management (RED)

Important:

On outbound traffic, policing and CBQ are mutually-exclusive features. To enable either feature, you must first ensure that the other is disabled.

In addition, monitoring is mutually-exclusive of policing or CBQ. If policing or CBQ is enabled, the enabled feature is automatically disabled when you enable monitoring. And if monitoring is enabled, it is automatically disabled when you enable policing or CBQ.

Assigning a policy map to interfaces

After you have configured a policy map, you can associate it with one or more interfaces. The policy map can be associated with an interface in the inbound or outbound direction. When you associate a policy map with an interface, an instance of the policy map is instantiated over the interface in the specified direction.

After you associate a policy map with an interface, you must enable the interface with the QoS features that are specified in the policy map. If you do not enable the specified QoS features, even though a policy map is assigned to the interface, no actions or classification are carried out.

Modifying a policy map

If you make any changes to a policy map that is already associated with one or more interfaces, then all applied changes are propagated to all the previously associated interfaces.

However, if the new configuration is not applicable to the previously associated interfaces, the changes are not applied to the interface. In this case, the new configuration is not applied to the mapped interfaces, but the policy remains mapped to the interfaces.

For example, if a CBQ configuration is applied to the inbound direction of a mapped interface, then the configuration is skipped for that interface.

But if the configuration is not valid, for example, exceeding the Policing CIR value for a particular mapped interface, then the configuration is not applied to any of the other mapped interfaces either.

Policy map clones

You can clone a policy map and save it to another policy map name to create a policy map that has minor differences from an existing configuration.

CoS marking

The CoS Marking feature in Chassis QoS sets the bits of well-defined fields in the data link layer or network layer header. This allows the QoS components of other routing devices to classify traffic based on the marked values. Although you can define your own values for marking fields for a packet, if you define the values to conform to standards, this provides greater uniformity in QoS handling of traffic flows between different networks.

The CoS Marking of a packet can be enabled at the following points in the system:

- Ingress classification
- Egress classification

You can classify any ingress or egress leaf classes with COS marking attributes. All the packets getting classified to the specified class are marked with the configured CoS attributes.

CoS marking supports marking of following attributes:

- IP precedence of IPv4 header
- User Priority (UP) of IEEE 802.1p Ethernet header
- DSCP of IPv4 or IPv6 header
- EXP value of MPLS header

IP Precedence marking

The IP precedence usually signifies the queue number on the egress port. However, in Chassis QOS, the queuing of packets is not based on IP precedence values. Instead, it is based on the CBQ priority value assigned by the user for the traffic class (see Queueing, scheduling, and shaping with CBQ on page 39).

The IETF recommended values for IP-precedence are listed below

Table 2: IP Precedence values

IP precedence	Type of application
7	Network control
6	Internetwork control
5	Critical
4	Flash override
3	Flash
2	Immediate
1	Priority
0	Routine

User priority marking

User Priority (UP) refers to the IEEE 802.1p User Priority field in the VLAN or priority-tagged packet header. The UP value ranges from 0 to 7 with 7 being the highest priority and 0 being the lowest.

Important:

With Q-in-Q stacked VLAN packets, user priority marking is performed only on the outer VLAN tag.

The recommended 802.1p values are listed in the following table:

Table 3: Recommended VLAN priorities for various types of traffic

802.1p (User priority)	Type of application
0	Best effort
1	Background and bulk transfers
2	Spare
3	Excellent effort
4	Controlled load
5	Video
6	Voice
7	Network Control Traffic

DSCP marking

DSCP refers to the DiffServ Code Point field in the IPv4 or IPv6 packet header. This value determines the Per Hop Behavior (PHB) given to the packet at each downstream node in the DS domain. The DSCP values range from 0 through 63. The mapping of DSCP values to different PHB is defined by a set of IETF standards, as shown in the following table:

Table 4: IETF recommended DSCP values

PHB	DSCP value	Binary
Default	0	000000
CS1	8	001000
AF11	10	001010
AF12	12	001100
AF13	14	001110
CS2	16	010000
AF21	18	010010
AF22	20	010100
AF23	22	010110
CS3	24	011000
AF31	26	011010
AF32	28	011100
AF33	30	011110
CS4	32	100000
AF41	34	100010
AF42	36	100100
AF43	38	100110
CS5	40	101000
EF	46	101110
CS6	48	110000
CS7	56	111000

The details of Assured Forwarding (AFxx), Expedited Forwarding (EF), and Class Selector (CSx) PHB values are described in RFC 2597 (AFxx), RFC 2598 (EF), and RFC 2474 (CSx)

DSCP remarking as policing action

Chassis QoS also supports the remarking of the DSCP value for a class as a policing action. You can configure a policing action for a leaf class to remark the DSCP value of the packet in the IP header based on its specified conformance level (see Traffic policing with srTCM and trTCM on page 28.)

Traffic policing with srTCM and trTCM

Traffic policing is one of the QoS components recommended in Differentiated Services compliant (DiffServ) traffic conditioning (RFC 2474 and RFC 2475). Policing meters traffic and performs actions based on the results of the metering. Metering compares the properties of the classified flows against a traffic profile and categorizes the traffic flows accordingly as conformed, exceeded, or violated. The traffic is then tagged as green, yellow, or red respectively.

The policer can perform actions against the packets based on the policing results. For example, packets that exceed the profile can be remarked in the DSCP field, and the packets that violate the profile can be dropped.

The color of the packet is specified by the DSCP value in the IP header. The conformance level of the packet, green, yellow or red is represented by a range of DSCP values. The following table gives the range of values for each conformance level.

Table 5: DSCP code point per Color

DSCP notation	Color
EF, CS6, CS7, AF11, AF21, AF31, AF41	Green
AF12, AF22, AF32, AF42, CS0-CS5	Yellow
AF13, AF23, AF33, AF43	Red

You can apply policing to any leaf classes that are identified by the Multifield traffic classifier.

You can enable policing in both inbound and outbound direction for all Chassis interfaces. In the inbound direction policing can be used for rate limiting the incoming flows. In the outbound direction, the output of the policing feature can be used in conjunction with the congestion management feature (see Congestion control and avoidance with RED on page 35).

Chassis QoS supports two different types of metering algorithms:

- srTCM
- trTCM

Both policing methods use a token bucket algorithm, where a token is considered as a byte and the bucket is a counter that is updated during the token refilling mechanism.

Policing using Single Rate Three Color Marker

srTCM (RFC 2697) uses a single rate Committed Information Rate (CIR) and two burst sizes, Committed Burst Size (CBS) and Excess Burst Size (EBS). The srTCM is implemented using two token buckets, namely Tc of size CBS and Te of size EBS. Both the token buckets are filled at the same rate (CIR).

When you configure srTCM, you must specify the CIR value. The CIR determines the token fill rate. You can also specify CBS and EBS but they are not mandatory values. By default CBS is one second worth of data corresponding to CIR and EBS is twice CBS. For example if CIR is configured as 10 Mbps then CBS is 10 Mb and EBS is 20 Mb.

A color of green is assigned by the meter if the packet does not exceed the CBS, yellow if it does exceed the CBS but not the EBS, and red if it exceeds the EBS. Coloring is only a means to convey the conformance level of packets. An action can be configured for each color. The supported actions include: Permit, Mark DSCP value, and Drop.

srTCM operation

The srTCM works as follows. Let the token buckets be represented by Tc(t) whose maximum capacity is CBS bytes, and Te(t) whose maximum capacity is EBS bytes. One token is the equivalent of one byte. Tc(0) is equal to CBS, and Te(0) is equal to EBS. Thereafter tokens are replenished as follows. If Tc < CBS then Tc is incremented by CIR bytes-per-second up to CBS; otherwise, if Te < EBS then Te is incremented by CIR bytes-per-second up to EBS.

When a packet of size B bytes arrives for policing, if Tc(t)-B >= 0, then the packet is marked green and Tc is decremented by B; otherwise, if Te(t)-B >= 0 then packet is marked yellow and Te is decremented by B else packet is marked red.

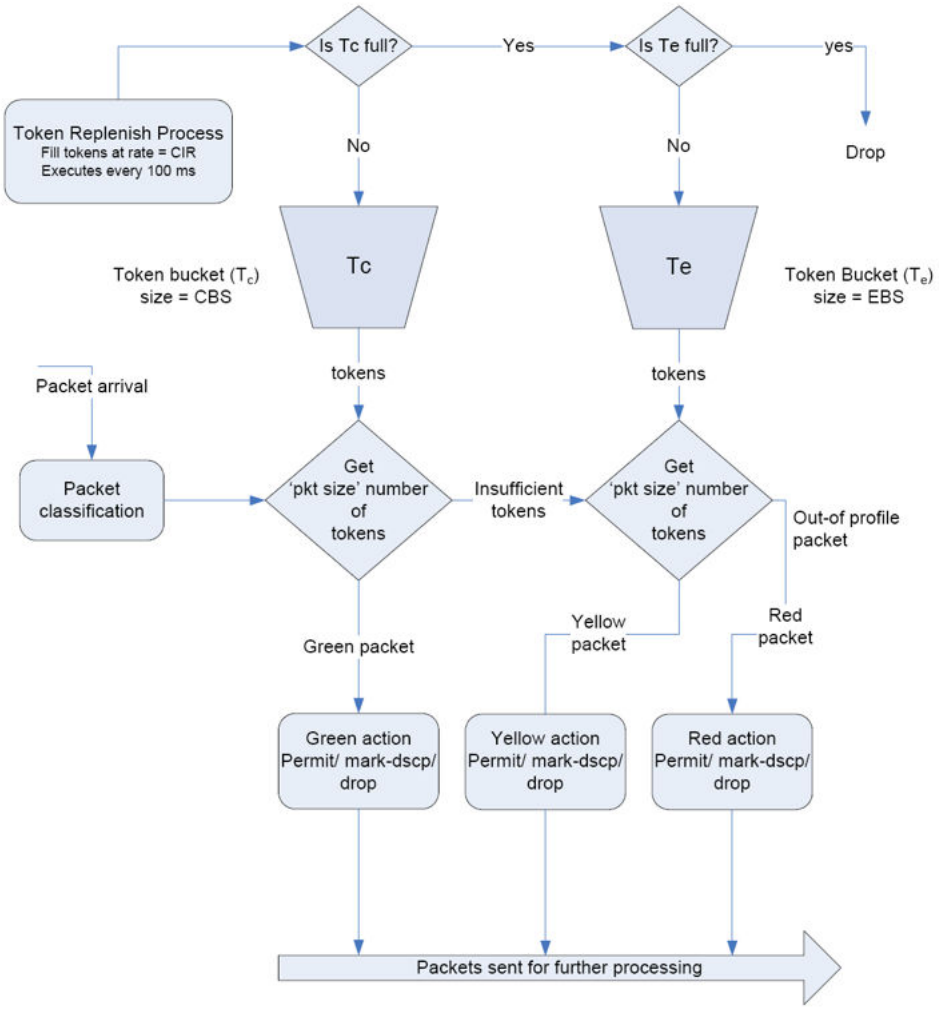

Figure 4: scTCM

The CIR is specified in Kbps. CBS and EBS are specified in Kbits. The Committed Burst Time (CBT) is defined by the relationship (CIR = CBS/CBT). The token filling process has to run at least every CBT milliseconds, without which CIR cannot be sustained.

Typically, it is recommended that CBS be large enough for CBT to be at least 1 second, in which case it is sufficient for the token filler to execute every 1 second and fill CBS number of tokens into the Tc and Te buckets.

However, with the Avaya Secure Router 2330/4134, the token filler is executed every 100 ms. This process replenishes tokens in all buckets configured on all interfaces in the system. Running the token filler more frequently results in more frequent and incremental updates to the token bucket, which can result in a smoother traffic pattern.

Policing using Two Rate Three Color Marker

The trTCM feature (RFC 2698) also utilizes two token buckets namely Tc of size Committed Burst Size (CBS) and Tp of size Peak Burst Size (PBS). But in this case, each token bucket has a different token fill rate, Committed Information Rate (CIR) and Peak Information Rate (PIR) respectively. You can configure the CIR, PIR, CBS and PBS.

A color of green is assigned by the meter if the packet does not exceed the CBS, yellow if it does exceed the CBS but is less than or equal to the PBS, and red if it exceeds the PBS. A policing action can be configured for each resulting color. Based on the conformance level, the associated action is performed on the packets. Supported actions include Permit, Mark-DSCP and Drop.

When you configure trTCM, you must specify the CIR and PIR values. You can also specify CBS and PBS but they are not mandatory values. By default, CBS is one second worth of data corresponding to CIR and PBS is one-second worth of data corresponding to PIR. For example, if CIR is configured as 10Mbps and PIR is configured as 20Mbps, then CBS is 10 Mb and EBS is 20 Mb.

trTCM operation

Let the token buckets be represented by Tc(t) whose maximum capacity is CBS bytes, and Tp(t) whose maximum capacity is PBS bytes. One token is the equivalent of one byte. The Tc(0) is equal to CBS, and Tp(0) is equal to PBS. Thereafter, tokens are replenished as follows. If Tc < CBS then Tc is incremented by CIR bytes-per-second up to CBS, and if Tp < PBS then Tp is incremented by PIR bytes-per-second up to PBS.

When packet of size B bytes arrives for policing, if Tp(t)-B < 0, then the packet is marked red else, if Tc(t)-B < 0 then the packet is marked yellow and Tp is decremented by B else the packet is marked green and both Tc and Tp are decremented by B.

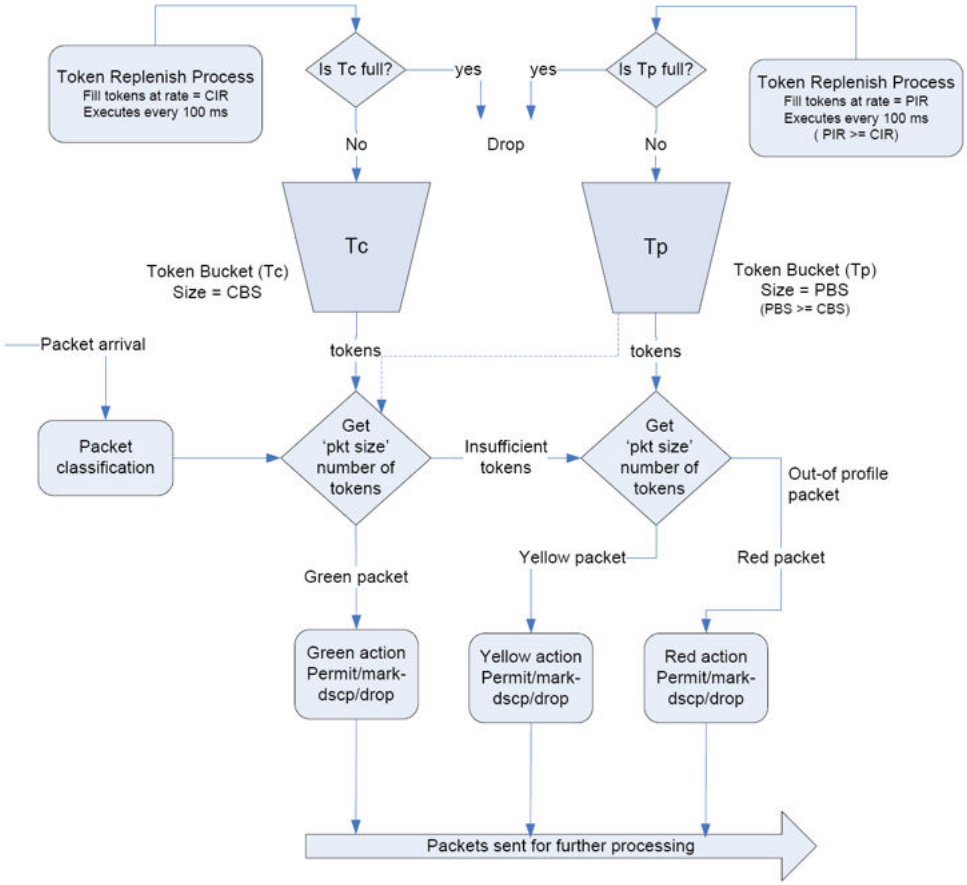

Figure 5: trTCM

The CIR and PIR is specified in Kbps. CBS and PBS are specified in Kbits. Committed Burst Time (CBT) and Peak Burst Time (PBT) are defined by the relationships (CIR = CBS/CBT) and (PIR = PBS/PBT). The period of the token filling process should at least be equal to the lesser of CBT and PBT milliseconds, without which CIR and PIR cannot be sustained.

Typically, it is recommended that CBS and PBS be large enough so that CBT and PBT are at least 1 second, in which case it is sufficient for the token filler to execute every 1 second and fill CBS number of tokens into the Tc bucket and PBS number of tokens into the Tp bucket.

However, with SR2330/4134, the token filler runs every 100 ms. Note that the same token-filling process replenishes tokens in all srTCM and trTCM token buckets. Running the token filler more frequently results in more frequent and incremental updates to the token buckets, which can result in a smoother traffic pattern.

If CBS and PBS are not specified, CBT and PBT are set to 1000 milliseconds and 2000 milliseconds respectively.

Policing actions

For each metering color, you can specify the following policing actions:

- allow the packet
- mark the packet with a specified DSCP value
- drop the packet

The default actions for each color are as follows:

- green: allow
- yellow: allow
- red: drop

With Chassis QoS based policing, you can choose any action for any result. It is not always necessary for a green packet to be allowed, a yellow packet to be remarked, or a red packet to be dropped. Any of the three actions can be assigned to any of the colors. This means that a green or yellow packet can even be configured to drop.

Color aware mode for srTCM and trTCM

Color aware mode is useful if the classified packet is already DiffServ conditioned, where the DSCP value carries the color information.

If classified packets are already DiffServ conditioned, you can enable the color-aware mode for srTCM or trTCM. In this case, the incoming packet is pre-marked with a DSCP value in the IP header as a result of a previous QoS treatment.

When srTCM or trTCM is enabled with color aware mode, they consider each packet as pre-colored packet before policing it.

The color aware mode can be enabled or disabled on per-class basis.

The incoming packet could be colored as green, yellow or red. The following table gives the action taken for each color when a colored packet arrives.

Table 6: Actions in color aware mode

Packet color	Conformance of packet	Policing action
Red	Either it conforms, exceeds or violates	Action defined for red packets is taken
Yellow	Conforms or exceeds conformance	Action defined for yellow packets is taken

Packet color	Conformance of packet	Policing action
	Violates conformance	Action defined for red packets is taken and remarked as red
Green	Conforms	Action defined for green packets is taken
	Exceeds conformance	Action defined for yellow packets is taken and remarked as yellow
	Violates conformance	Action defined for red packets is taken and remarked as red

The advantage of using color aware mode is that this reduces the steps of policing algorithm.

The policing feature enables Chassis QoS to police the traffic corresponding to micro flow, aggregate flow, or behavior flow. For example, VLAN ID based policing can be enabled on ingress traffic by configuring a traffic class with VLAN ID as the classification key and setting the appropriate policing parameters for the traffic class. Similarly for application based policing, configuring traffic class with the application identifier as a classification key and configuring the policing parameters for traffic class can achieve policing.

The policing color information can be used in other traffic conditioning functions such as congestion management to trigger a particular action for each packet, which is either in-profile or out-of-profile.

Policing versus shaping

Policing differs from shaping, which controls the traffic by using a queuing mechanism to delay packets that arrive faster than the configured rate. The shaping can smooth out the burstiness in a flow, helping to minimize buffer overruns in intermediate routers.

Policing does not smooth out bursts, but it can control bursts by dropping packets or by reducing the priority of the packets by remarking the DSCP field of the IP header of the packet. Since shaping delays packets instead of dropping them, it is more suited to adaptive applications (those using TCP). With policing, dropping packets can cause exponential back off of TCP, which can adversely affect the throughput. However, this behavior can be minimized by setting the policing parameters to a sufficiently large value.

Policing has the advantage of providing low latency since it does not queue packets. This makes policing a good choice for interactive and streaming voice and video applications. Policing also uses fewer resources in the router than shaping. It is a better method for providing QoS for incoming traffic on an interface.

Traffic monitoring

Traffic monitoring is a QoS feature used for measuring the traffic flow. If an interface is enabled with the QoS traffic monitoring feature, the packets are classified and the statistics such as counters and bandwidth usage are updated. After you have configured the desired traffic classes using the appropriate match rules, you can enable Traffic Monitoring to measure the flows.

For any interface, only one of the following QoS features can be enabled at any one time: CBQ (scheduling and shaping), Policing, and Monitoring. These three features are mutually exclusive. Therefore, when traffic monitoring is enabled, CBQ and policing are automatically disabled. While monitoring is running, Chassis QoS does not perform any other treatment for the classified packet, except marking and policy-based redirect if configured. The packets get classified into traffic classes and the class statistics are updated without applying any QoS treatment.

This feature is useful to understand the bandwidth consumption of different traffic flows on the network, which in turn can aid in the capacity planning of the network and in defining QoS policies.

Traffic Monitoring can be enabled for ingress and egress traffic of any interface.

By default the monitoring feature is disabled.

Congestion control and avoidance with RED

With Chassis QoS, you can enable RED (Random Early Detect) to perform congestion avoidance on output queues.

The basic operating philosophy of RED is that it detects the onset of congestion and starts dropping packets in a random fashion before queue overflow leads to tail drops. Random drops not only improve throughput of adaptive applications using TCP but are also more suitable than tail drops for voice and streaming audio or video, as they result in less perceptible degradation.

RED supports three levels of drop precedence (DP) for handling congestion on egress queues. When congestion occurs on an egress queue, packets with higher DP can be configured to drop, while packets with low DP are queued

The three levels of Drop Precedence are applied to the queue based on the color (red, yellow, green) values of DSCP of the packet.

The SR2330/4134 supports two types of RED: WRED (Weighted Random Early Detect) and DS-RED (Differentiated Service Random Early Detect). WRED and DS-RED perform the same

basic operations, but with WRED, you assign RED DP parameters to queues, while with DS-RED, you can assign RED DP parameters to each DSCP code point in a queue.

The following sections define the common characteristics of WRED and DS-RED.

Average queue size

When a packet arrives in a queue, RED calculates the average queue size and not the current queue size to determine congestion. The average queue size is calculated as the EWMA (Exponential Weighted Moving Average) of the current queue size. Simply put, the average queue size is a smoother version of the instantaneous queue size.

RED uses this average queue size as an indication of congestion, and compares it against the minimum threshold (minth) and maximum threshold (maxth) to determine the appropriate action for the packet:

- If the average queue size is less than the minimum queue threshold (minth), the arriving packet is queued.
- If the average queue size is between the minimum queue threshold and the maximum threshold (maxth), the packet is either dropped or queued, depending on the packet drop probability.
- If the average queue size is greater than the maximum threshold, the packet is automatically dropped.

The average queue size is based on the average calculated for the last received packet, and the current size of the queue. The average queue size is calculated using the following formula:

Qaverage = (1 – w) * Qaverage + w * Qcurrent Where w = $1/2^{ewf}$ and EWF = Exponential Weighing Factor

EWF

The Exponential Weighing Factor (EWF) is a user-configurable parameter that allows you to control the average queue size. The valid values for EWF range from 1 to 15.

For high values of EWF, the previous queue average becomes more important. A large weighting factor flattens the peaks and lows in queue length. The average queue size is unlikely to change very quickly, avoiding drastic swings in size. The RED process is slow to start dropping packets, but it can continue dropping packets for a time after the actual queue size has fallen below the minimum threshold. The slow moving average accommodates temporary bursts in traffic. However, if the value of EWF gets too high, RED cannot react to congestion. In which case, packets are transmitted or dropped as if RED was not in effect.

For low values of EWF, the average queue size closely tracks the current queue size. The resulting average may fluctuate with changes in the traffic levels. In this case, the RED process responds quickly to long queues. Once the queue falls below the minimum threshold, the

process stops dropping packets. However, if the value of EWF gets too low, RED overreacts to temporary traffic bursts and drop traffic unnecessarily.

Packet drop probability

The packet drop probability is based on the minimum threshold (minth), maximum threshold (maxth), and mark probability denominator (mpd).

When the average queue size is above the minimum threshold, RED starts dropping packets.

The rate of packet drops increases linearly as the average queue size increases. The rate of packet drops moves linearly from zero (when average queue size is below minth) to the mpd value (when average queue size is equal to, but not exceeding, maxth). Once the average queue size exceeds maxth, then all packets are dropped.

The mpd value is the fraction of packets dropped when the average queue size is equal to, but not above, the maximum threshold. For example, if the mark probability denominator is 512, one out of every 512 packets is dropped when the average queue is at the maximum threshold.

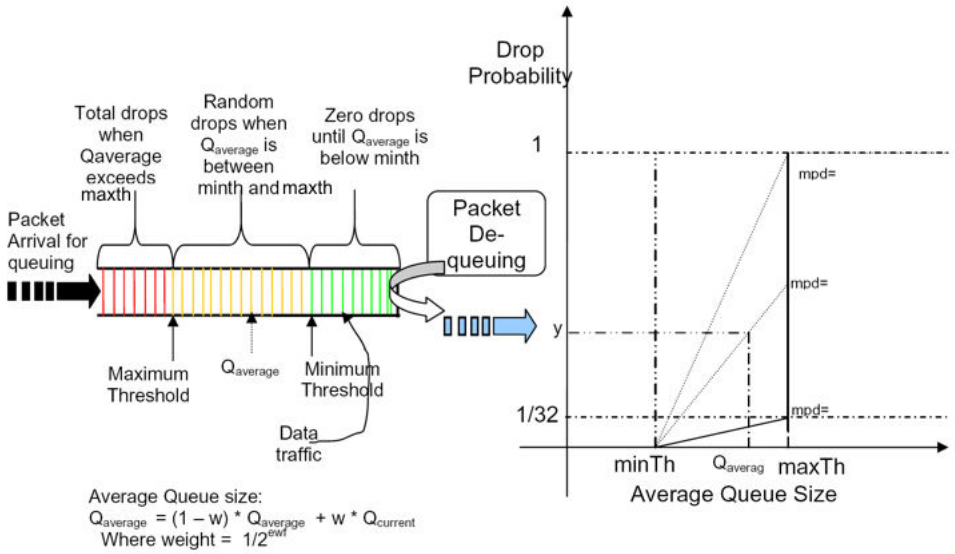

Figure 6: Packet drop probability

Set the minimum threshold value high enough to maximize the link utilization. If you set the minimum threshold too low, packets can be dropped unnecessarily, and the transmission link is not fully used.

Set the difference between the maximum threshold and the minimum threshold to be large enough to avoid global synchronization. If the difference is too small, many packets can be dropped at once, resulting in global synchronization.

The threshold values are dependent on the bandwidth of classes (and in turn the bandwidth of the interface). If the bandwidth of the interface or class changes, the threshold values are adjusted to the newer bandwidth.

The following table provides the default percentage values (percentage of queue size) for each color along with the range of DSCP values for each color.

Table 7: Congestion threshold (as recommended by NSC)

DSCP notation	Color	RED minth and maxth values
EF, CS6, CS7, AF11, AF21, AF31, AF41	Green	75%-99%
AF12, AF22, AF32, AF42, CS0-CS5	Yellow	70%-90%
AF13, AF23, AF33, AF43	Red	40%-70%

Handling of control packets

The RED algorithm for congestion management does not drop the control packets, unless the average queue size exceeds the configured maximum queue size or there are insufficient resources like buffers. In case of shaping, dropping of control packets is due to the average queue size being greater than allowed maximum queue size or due to insufficient resources. In other words, the control packets are dropped only when the queue overflows due to congestion in the queue, which is a rare occurrence because RED avoids congestion.

DS-RED

While WRED allows you to specify minth and maxth thresholds for the three DSCP colors only, DS-RED allows you to specify these thresholds separately for each DSCP value. The values for maxth and minth thresholds as well as ewf and mpd can be different for each DSCP value.

When RED is enabled

You can apply RED in the outbound direction only. RED is applied to the QoS flow in the following cases:

- If no QoS features are enabled on an interface, or if you enable policing on an interface, WRED is automatically enabled at the interface level. In these cases, you can also enable DS-RED at the interface level.
- If you enable CBQ on an interface, then WRED is automatically enabled for each of the specified classes. (In this case, interface-level RED is not applied to traffic.)

  You can also enable DS-RED at the class level, provided that the classification type for the configured leaf class is DSCP.

WRED and DS-RED are mutually-exclusive at the interface and class level. Both features cannot be enabled on the same interface or class.

Queueing, scheduling, and shaping with CBQ

Class Based Queuing (CBQ), also called Flow Based Queuing, is a queuing, scheduling and shaping mechanism used to deal with output congestion on applicable interfaces by intelligently prioritizing the traffic and managing the available bandwidth. With CBQ enabled, limited resources are put to best use during times of congestion.

CBQ can manage the bandwidth for all outbound traffic on an interface. When the CBQ feature is enabled on an applicable interface, all outbound traffic on that interface is classified, queued, scheduled, and shaped by the CBQ mechanism.

The CBQ feature can only be enabled on the outbound direction for chassis interfaces. No traffic shaping is done for inbound traffic.

At the heart of CBQ is a CBQ scheduler. The CBQ scheduling algorithm uses Periodic Deficit Round Robin with Priority (DRRP-P). The algorithm encompasses features of both Weighted Fair Queuing (WFQ) and Deficit Round Robin (DRR). It also supports the setting of Priority and Peak Rate limiting for traffic flows.

The following figure depicts the typical set of operations involved in CBQ operation.

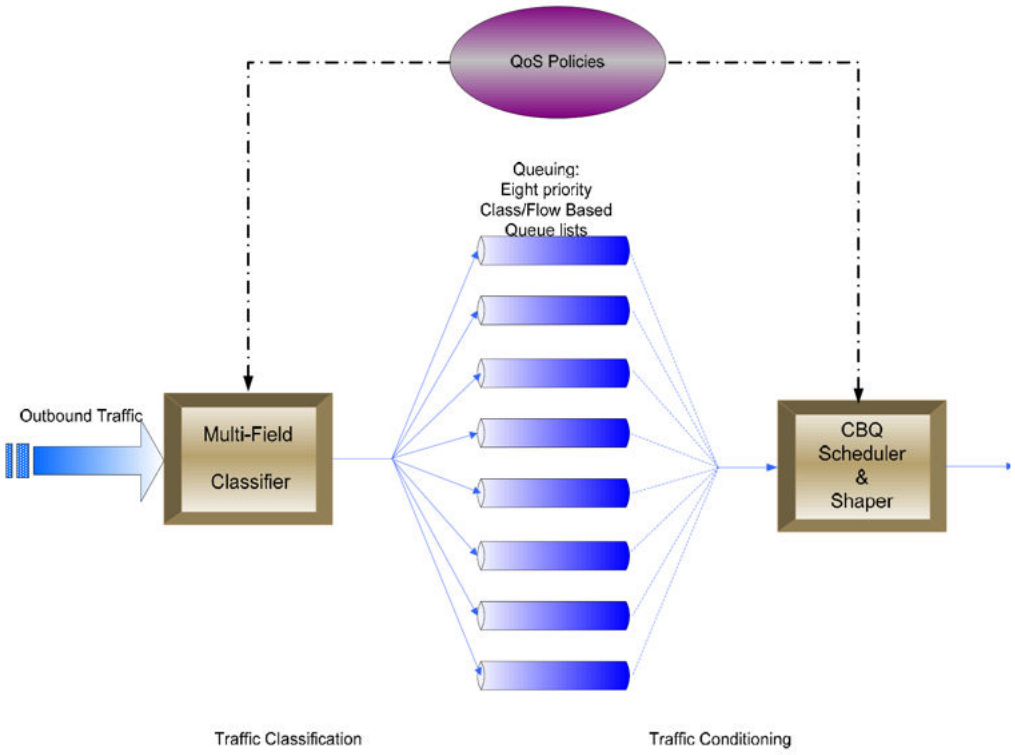

Figure 7: CBQ operation

In the preceding figure, the outbound traffic is comprised of forwarded packets from the inbound direction of interfaces and locally generated traffic. The inbound and locally generated traffic transitions to outbound traffic during forwarding and routing. Then the outbound traffic is fed to the QoS engine.

In the QoS engine, the outbound traffic first passes through the multifield classifier, and then through the CBQ queuing, scheduling and shaping stages.

Queuing and priority

In the CBQ mechanism, each interface has eight priority queue lists. Each priority queue list is comprised of a list of queues. These queues are represented by the leaf classes in the outbound hierarchical classification tree.

The priority queue lists prioritize the order in which traffic flows are serviced. Each leaf class or queue in the outbound hierarchical classification tree can be assigned to one of the 8 priority queue lists.

Priority queue list 1 has the highest priority, and priority queue list 8 has the lowest. Leaf classes in the higher priority queue lists are serviced before lower priority ones and therefore

experience lower latency. Leaf classes of the same priority experience approximately the same average queuing latency and excess bandwidth is also equally shared between them.

To associate a leaf class with a particular priority queue list, you can configure the priority parameter.

There are no limitations or restrictions on the number of leaf classes per priority queue list. However, the greater the number of leaf classes, the higher the impact is on performance and memory consumption.

All queues in the priority queue lists are serviced by the scheduler at least once in every 5 msec.

To specify the level of service the leaf class gets from the scheduler, you can configure the following CBQ parameters: Committed Rate (CR), Peak Rate (PR), and Priority. Regardless of the configured priority value, a leaf class is always guaranteed the configured CR of bandwidth.

Committed Rate for bandwidth guarantee

Each leaf class can be guaranteed a minimum amount of bandwidth at all times, even during congestion. The Committed Rate (CR) parameter is used to specify the amount of guaranteed bandwidth for a leaf class. The CR is specified in terms of a percentage of the interface bandwidth. The sum of the CR for all leaf classes on an interface cannot exceed the total bandwidth of the interface.

The CR is a mandatory parameter. Without it, the queue for the leaf class is not added to the priority queue list. The other parameters can assume their default values if not specified. The default value for PR is CR, and for priority it is 8.

Bandwidth sharing

If a leaf class needs bandwidth above the configured CR, it can borrow any excess bandwidth available on the interface. The excess bandwidth is available when one or more of the other leaf classes is using less than its CR or when the sum of CRs of leaf classes configured on the interface is less than the interface bandwidth.

Classes are allowed to borrow excess bandwidth in the order of their priority. A higher priority class has preferential access to any excess bandwidth over a lower priority class.

However, the bandwidth offered to a leaf class never exceeds the configured PR for that leaf class.

Peak Rate limiting

The Peak Rate (PR) parameter is used to limit the bandwidth consumption of a class to a maximum value. The peak rate is specified in terms of a percentage of interface bandwidth.

Peak Rate limiting for child nodes using parent nodes

As shaping happens only for successfully classified flows, the CR and priority parameters are meaningful only for leaf classes. However, the PR parameter is meaningful for all classes. In the classification hierarchy, if you configure the PR for a parent node, all child classes under that node are collectively peak-rate-limited to that value. This means that the sum of the CR for all child classes cannot exceed the PR value of the parent class.

DRRP-P scheduling

Shaping is used to rate-limit a flow by queuing packets and subsequently scheduling the queue at the configured shaping rate. Unlike policing, shaping smoothes out the burstiness in a flow. CBQ can be used to shape a flow by making use of the PR (Peak Rate) parameter. PR should be set to the desired shaping rate.

The Periodic Deficit Round Robin with Priority (DRRP-P), like DRR, the algorithm accounts for packet sizes by keeping track of the deficit for each class. In addition it supports priorities and is periodic in nature.

The DRR-P scheduling algorithm performs the following:

- guarantees bandwidth for each class/queue.
- offers bandwidth based on bytes and not packets. Therefore it does not allow flows with bigger packet sizes to dominate.
- runs periodically every 5 msec. This limits the queuing delay experienced by each queue/class because each class is serviced at least once every 5 ms.
- supports prioritization in terms of both latency and bandwidth borrowing.
- supports rate limiting of classes to a configured maximum value.

Strict Priority Queuing (SPQ)

With SPQ, the scheduler always services higher priority queues before lower priority ones. There is no bandwidth configuration involved. The scheduler services a given queue only if there are no packets waiting in a higher priority queue. This ensures that the highest priority queues are always serviced - although it can lead to starvation of the lower priority queues.

To configure the CBQ scheduler to behave as a Strict Priority Scheduler, assign the desired priority to each queue, but configure all queues with a CR of 1 and PR of 100 (full interface bandwidth). In this case, the CBQ scheduler services queues only on the basis of their priority.

Note that, for SPQ, you must set the CR to 1 as the value cannot be set to 0.

Congestion avoidance

The following sections describe congestion avoidance in the context of CBQ.

WRED on each class queue

Weighted RED is enabled by default on each class queue for congestion avoidance. Weighted RED thresholds can be individually configured for each class queue. (see Congestion control and avoidance with RED on page 35)

DS-RED for each DSCP code point

If a classification type for a leaf class is DSCP, DS-RED can be manually enabled on a per class queue for congestion avoidance. DS-RED thresholds can be individually configured for each DSCP code point which is assigned to the class. (see DS-RED on page 38)

Congestion avoidance disabled

When an outbound packet is classified to a leaf class, the packet is queued into that class queue after subjecting to queue admission control, which is administered by the congestion avoidance algorithms.

If congestion avoidance is disabled for leaf classes, packets are queued in to the class packet queue as long as the current queue size is less than maximum allowed queue size. Otherwise packets are dropped and increments the buffer over flow counters. The optimum value for maximum allowed queue size for class packet queue is by default determined by the CR and PR parameters, and additionally you have the ability to modify it.

PVC behavior when CBQ enabled

On Frame Relay bundles, whenever a PVC is created, it uses all the available or unassigned bandwidth. When CBQ is enabled either manually or using auto QoS, subsequent PVC creation fails because there is no available bandwidth. To avoid this scenario, you must change the bandwidth for created PVCs such that there is always non-zero bandwidth available.

QOS Strict Priority Queuing

Secure Router implements Strict Priority Queuing (SPQ) to minimize latency and jitter for traffic over Ethernet and Bundle interfaces. SPQ uses the shaping/scheduling infrastructure currently used with Class Based Queuing (CBQ), so there is minimal change to QoS configuration. Traffic is classified and marked as before using Policy-maps. Each traffic flow (class-map) is mapped to a priority queue. Multiple flows can be mapped to the same priority queue. When SPQ is enabled, instead of queuing the classified traffic into class queues, the traffic will flow through one of the interface queues based on the configuration. SPQ supports up to 8 queues per interface with pre-defined priorities. Queue 1 is highest priority queue while queue 8 is the lowest priority queue. All unclassified traffic is placed in queue 8 by default. SPQ can be enabled or disabled at the interface level for outbound flows. Only CBQ or SPQ can be active on any interface yet both can be active at the same time on different interfaces. SPQ is only supported on Ethernet, PPP, FR, MLPPP & MFR interfaces.

Unlike CBQ, where the committed rate percentage and peak rate percentage are specified globally in the class map, with SPQ, committed rate percentage is specified for each queue at the interface level with the shape command. The committed rate percentage can be configured between 0% (default) and 100% of the interface bandwidth. The peak rate percentage is 100% for all priority queues and cannot be modified. If the queues are configured with committed rates, they are serviced in round-robin mode. Any bandwidth available after all the committed rates are fulfilled is used to service the queues in strict-priority mode. Also, WRED can be configured on each queue for congestion control on Bundle interfaces. The latency for traffic for a SPQ queue can increase once it exceeds the committed rate percentage for that queue.

QOS SPQ CLI Commands

To configure SPQ, you must first setup up the policy map under the qos chassis section. A policy map consist of class maps where each class map is assigned a SPQ queue number with the assign-queue command. Shaping is configured for the queues using the shape command under interace queue section. When SPQ is enabled on an interface, the committed rate and peak rate percentages defined in the class map are ignored. All the clear and show commands are equivalent on SPQ as for CBQ.

Policy-based redirect

With policy-based redirect, you can specify that packets from a specified class are to be forwarded to a destination specified by any one of the following parameters:

- IPv4 address of the next hop
- IPv6 address of the next hop

- Interface name (bundle name or tunnel name)
- MPLS LSP name

Policy-based redirect is applied at the ingress.

Buffer Management

The Buffer Management (BM) module manages the global common buffer pool (or stack) which is shared by all interfaces in the system. In this case, the buffers being discussed refer to data buffers used to hold received packets. Buffers for the global buffer pool are pre-allocated at system start-up. The total number includes buffers for attaching to the LAN and WAN descriptors, "queue buffers" for transmit queues on WAN interfaces, and buffers needed for reassembly on WAN receive. .

In most cases, you do not need to configure buffer management. When a new class map is added or deleted or when CR and BR of a class map are changed, appropriate buffer management changes are automatically made.

The number of "queue buffers" in the global pool is proportional to the bandwidth of applicable interfaces (for those interfaces which support the CBQ feature) that the system supports which is indicated by the model number.

When an applicable interface is created, a certain amount of queue buffers are reserved for its transmitting queue from the global buffer pool. Buffer reservation is only done for applicable interfaces on which we expect congestion and hence queue build-ups on those applicable interfaces. The number of buffers reserved for an applicable interface’s queue is such that it provides 15 msec of buffering for 100 byte packets at a rate equal to the interface bandwidth. This amount of buffering is always guaranteed for an applicable interface. An applicable interface can use more than its reserved amount of buffers by borrowing queue buffers from the global pool if available. However, only the portion allocated for transmit queues can be borrowed. Therefore it is not possible for traffic, flooding an applicable interface, to consume all buffers and starve for receive buffers on other interfaces.

When CBQ is enabled on an interface, many more buffers are needed because buffers have to be reserved for each queue. In this case, the number of queue buffers reserved for the interface is such that traffic at the full rate of the interface can buffer for 50 msec, considering 100 byte packets. For a single T1 PPP interface with CBQ enabled, the number of reserved buffers would be 96 compared to 29 when CBQ is disabled.

You cannot change the number of reserved buffers for an interface using the CLI. When the bandwidth of an applicable interface is changed administratively, the number of reserved buffers is automatically updated. If the total number of queue buffers to be reserved for all interfaces in the system increases beyond what is preallocated in the global pool, more buffers are dynamically allocated and added to the global pool. This can happen when many interfaces in the system are enabled for CBQ. Later, when buffers are returned to the common pool, because some interfaces were deleted or CBQ was disabled on them, and the number of queue buffers in the pool becomes more than the original preallocated value, the excess is

freed back, thereby making more system memory available for other purposes like routing tables.

When class maps are created on applicable interfaces and CBQ is enabled, a certain minimum number of buffers are reserved for each class map from the interface's buffer pool. The remaining buffers (which are not reserved) form a "common class map buffer pool" which can be shared among all class maps of the interface. When a class map has used up all its reserved buffers (for queuing packets), it can loan buffers from the "common class map buffer pool" but only up to a maximum limit configured for the class map. This "Max Buffer" limit for a class map is selected in such a way that no single class map can consume all the buffers in the common pool. Given this fact and the fact that all class maps have a certain minimum reserved buffers, it is not possible for malicious traffic flow in one class map to starve other class maps out of buffers and deny them service.

bbr is the maximum bytes that a class map can send and bcr is the guaranteed number of bytes that a class map can always send in one scheduler interval (which is 5 ms). Therefore, in order for the class map to have a bandwidth of PR Kbps, the class map queue should have enough buffers to hold at least bbr bytes in each interval. In our buffer calculation, as always, we will consider the average packet size to be a low value of 100 bytes which will guarantee enough buffers for typical internet traffic. We reserve (2*bbr/100) or (4*bcr/100) buffers for each class map, whichever is higher. However, it is possible for many or even all class maps on an interface to have very small CRs but BR equal to the interface bandwidth. This will significantly increase the total buffer requirement for the interface. We therefore limit the maximum buffers to be reserved for a class map to (6*bcr/100). It is possible for a class map to have a very low CR (say 10Kbps). In this case we reserve at least 4 buffers for the class map. The complete equation to calculate the reserved buffers for a class map is given below along with example.

Since at least 4 buffers need to be reserved for each class map, the maximum class maps that can be configured on an interface are (buffers in interface pool / 4). For a 1 T1 bundle the maximum class maps that can be created is 96/4 = 24. For a 2 T1 bundle, it would be approximately double.

The "Max Buffers" for a class map is configurable on the CLI. The default and recommended value is (3*bbr/100) or (10*bcr/100) or (10) whichever is the maximum. The equation along with an example is given below.

You can modify the "Max Buffers" limit for a class map. By lowering this limit, a class map can be restricted to using fewer buffers which in turn increases the chances of buffer availability in the common pool for other class maps. The upper bound of this value depends on the number of buffers available in the common class map pool and the lower bound depends on the maximum of max threshold value used in congestion avoidance algorithm, reserved buffers to offer CR/PR, and minimum required buffers to offer PR rate. The valid range is shown on the CLI class map configuration display.

Since, (sum of bcr of all class maps) <= Bc, in a typical configuration, the total number of buffers reserved for all class maps will be not more than (6 * Bc)/100. Therefore, there will be at least (4 * Bc)/100 buffers in the common class map pool.

QoS over Frame Relay

The Frame Relay protocol provides basic in-built QoS functionality such as PVC shaping, and congestion handling through Forward Explicit Congestion Notification (FECN) and Backward Explicit Congestion Notification (BECN) fields. However, these Frame Relay mechanisms alone do not provide QoS services such as Differentiated Services (DiffServ) compliant behavior, or prioritization of the traffic.

You can use QoS over Frame Relay to enable all the same features that are implemented on any other chassis interfaces, like flow classification, flow-based policing, monitoring, marking, queuing, and shaping. The functional behavior of these features remains the same.

The major difference when QoS is applied to Frame Relay is that the policy map configuration is applied at the logical interface level. In other words, rather than applying the policy map to an interface, you must apply the policy map to the virtual circuit (PVC). Although you must still enable the QoS features at the interface (bundle) level.

The other behavior to note when you configure QoS on Frame Relay relates to CBQ. When you configure CBQ for a Frame Relay PVC, the percentages represented by the committed rate are relative to the PVC speed, while the percentages represented by the peak rate are relative to the total bundle speed. This allows a PVC, or a class within a PVC, to burst to the bandwidth of the bundle.

And if all PVCs are equally loaded, each PVC gets an equal amount of bandwidth, which is approximately equal to the total bundle speed divided by number of PVCs in the bundle.

The sum of CR of all leaf classes in a bundle must not exceed 100%.

MPLS QoS

MPLS QoS provides support for global DSCP-to-EXP mapping on the ingress LER, and global EXP-to-DSCP mapping on the egress LER. On the ingress LER, MPLS QoS also supports flow-based EXP marking for inbound traffic, and class-based queueing for outbound traffic.

The following sections provide more details on these features.

Ingress LER- EXP marking

In order to give fair and expected QoS treatment for various traffic flows funneling through the MPLS LSP tunnels, each packet must be marked with the correct EXP value on the ingress

LER. The following are the available methods of mapping/marking of the EXP value for packets on the ingress LER:

- Global DSCP-to-EXP Mapping
- Flow-based EXP Marking

If provisioned, these methods can operate in tandem.

MPLS QoS also supports class based queuing of per-EXP traffic on the ingress LER.

Global DSCP-to-EXP Mapping

In the ingress QoS processing stage of ingress LER, by default, every packet is marked with the EXP value based on the global DSCP-to-EXP mapping table shown below. For any packet, if DSCP is not applicable, then the EXP value corresponding to the DSCP value of 0 is marked.

Each MPLS per-EXP flow is serviced at the defined priority and bandwidth. The peak rate allows LSP flows to utilize the unused bandwidth up to the full interface bandwidth.

Table 8: Global DSCP to EXP mapping

Class	DSCP	EXP	Bandwidth allocated per EXP within LSP (specified as % of LSP, unless otherwise stated)
Critical Control Traffic	Class Selector 7	7	CR: 10%, PR: 100% of interface Tail Drop, Priority : 1
Network Control Traffic	Class Selector 6	6	CR: 10%, PR: 100% of interface Tail Drop: Priority: 2
Real Time	EF	5	CR= 35%, PR=50%, Tail Drop, Priority: 3
Class 1	AF 4X	4	CR=10%, PR: 100% of interface, Priority: 6
Class 2	AF 3X	3	CR=10%, PR: 100% of interface
Class 3	AF 2X	2	CR=5%, PR: 100% of interface, Priority: 6
Class 4	AF 1X	1	CR=10%, PR: 100% of interface Priority: 7
Best Effort	Default	0	CR=10%, PR: 100% of interface Priority: 8

Comments? infodev@avaya.com

Flow-based EXP Marking

Flow-based EXP marking is supported on the inbound direction only. You can use multifield classification to define to define traffic classes, and specify the desired EXP marking as the action on leaf classes.

Class-based queueing

MPLS QoS also supports class based queuing of per-EXP traffic, based on the EXP value of the data after applying the global DSCP-to-EXP mapping, and flow-based EXP marking, if applicable.

DSCP Marking on Egress LER

In order to give fair and expected QoS treatment for various traffic flows coming out of the MPLS LSP, each of the packets can remarked with proper DSCP code points in the egress LER. The following are the available methods of marking the DSCP code points for packets on the egress LER.

- Global EXP-to-DSCP marking
- Flow-based DSCP marking

If provisioned, these methods can operate in tandem.

Global EXP-to-DSCP Marking

In the ingress QoS processing stage of egress LER, by default, every packet is re-marked with the DSCP value based on the global EXP-to-DSCP mapping table.

In the egress LER, the changes to the EXP values along the MPLS network path can be reflected in to the packet by appropriately re-marking the DSCP value.

The following table provides the default EXP-to-DSCP mapping per EXP class in the egress LER.

Table 9: Global EXP to DSCP mapping

Class	EXP	DSCP
Critical Control Traffic	7	Class Selector 7
Network Control Traffic	6	Class Selector 6
Premium, Real time	5	EF
Platinum, Class 1	4	AF 41

Class	EXP	DSCP
Gold, Class 2	3	AF 31
Silver, Class 3	2	AF 21
Bronze, Class 4	1	AF 11
Best Effort	0	Class Selector 0, Default

The EXP-to-DSCP functionality is dependent on the configured MPLS tunnel mode. The tunnel modes control whether the DiffServ markings for IP packets at the egress LER remain independent from, or are a function of, the MPLS label EXP values. These modes have no influence on intermediate LSRs.

There are three tunnel modes that control the application of EXP values in various scenarios:

Uniform mode

In uniform mode, changes made to the EXP value on the uppermost label are applied to all labels in the stack, including the IP packet.

In the egress LER, the changes to the EXP values along the MPLS network path are reflected into the packet by appropriately re-marking the DSCP value based on the global EXP-to-DSCP mapping table.

Pipe mode

In pipe mode, changes made to the EXP value on the uppermost label are propagated to other MPLS labels but not to the IP packet. In this case, the DSCP value of the IP packet does not change from ingress LER to egress LER. However, the PHB at the egress LER is chosen based on the removed EXP value.

Short-pipe mode

Changes made to the EXP value on the uppermost label are propagated to other MPLS labels but not to the IP packet. In this case, the DSCP value in the IP packet remains unchanged. At the egress LER, the PHB is not chosen based on the removed EXP value, but rather on the DSCP value of the IP packet.

Flow-based DSCP Marking

Flow-based DSCP marking is supported on inbound or outbound direction of the egress LER. You can use multifield classification to define traffic classes and assign DSCP marking as an action.

Crypto QoS (CBQ) for IPSec VPN

The SR2330/4134 crypto engine preforms encryption and hashing of packets for IPSec VPN tunnels. However, the crypto engine throughput is less than the system throughput, and therefore, there is potential for congestion to build up and packets to be dropped at the crypto

engine queue. You can apply CBQ to packets entering the encryption engine to guarantee bandwidth and reduce the latency for delay sensitive voice packets.

As shown in the following figure, the crypto engine maintains an input FIFO queue and an output FIFO queue. The packets needing crypto services (encryption, decryption, hashing, and so on) have to be queued into the input FIFO queue for the crypto engine to act on it. When the processing is completed, the packet is placed on the crypto output queue.

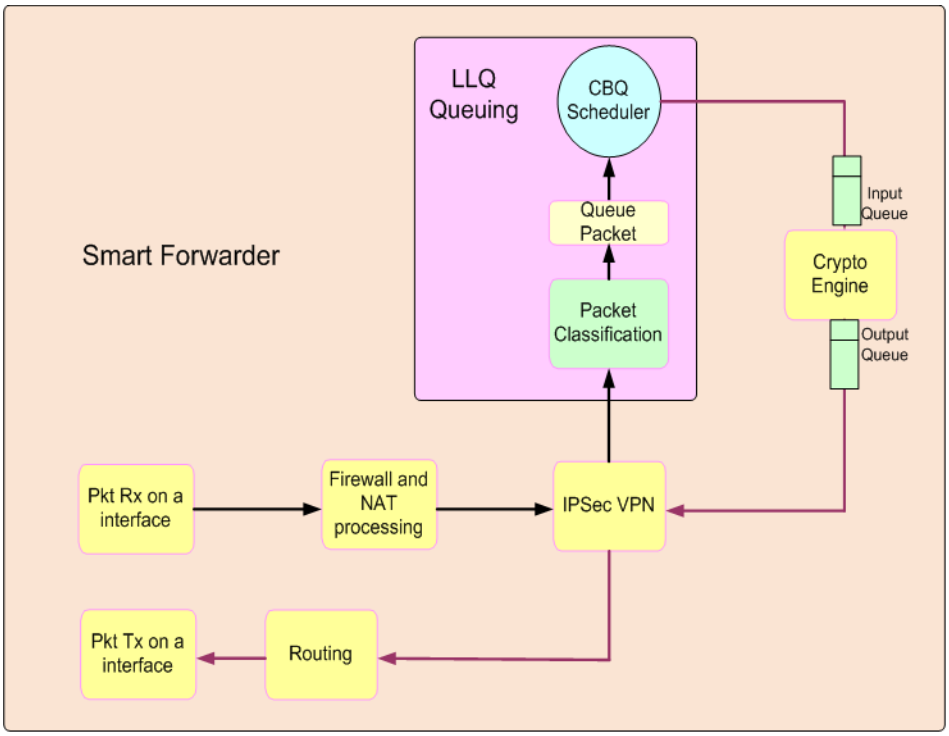

Figure 8: LLQ

To avoid congestion on the crypto engine, you can classify traffic into the desired classes using the QoS multifield classifier and apply CBQ with associated committed rate and priority parameters for each crypto class.

Crypto interfaces do not support RED or policing, although marking can be configured.

Packets that are classified into a leaf traffic class are placed in the associated class queue. The CBQ scheduler periodically services the queues of all leaf traffic classes on each crypto interface. The service that each class queue receives depends on the service parameters (CR, PR, Priority) assigned to it. Eight priority levels are supported, from priority 1 (highest) to priority 8 (lowest).

The crypto engine bandwidth varies for different packet sizes. Therefore, for each class you must configure the committed rate as a percentage of the crypto interface bandwidth. The average bandwidth for the interface is calculated by the number of bytes processed by the crypto engine in a given interval. This bandwidth is used to calculate the average interface bandwidth using the exponentially weighted moving average. This ensures low latency without sacrificing the interface bandwidth.

Historical statistics

When traffic classes are created and QoS features are enabled, per-class traffic statistics are collected during traffic classification. You can view these statistics on the system console using CLI commands. But with the CLI, you can view only the current system statistics.

With the historical statistics feature, you can upload historical QoS statistics to an FTP server at configurable intervals (1, 2, 3, or 4 hours). You can also configure the sample interval for the statistics (5, 10, or 15 minutes). During the upload interval, statistics are collected at every sample interval.

The statistics are saved in a tab delimited text file. For each upload, a new file is created in the FTP server, with the upload time indicated in the file name.

Collection of historical statistics is always enabled. It cannot be disabled. However you can enable and disable uploading to an FTP server. You can also specify the name of the file to be uploaded. The system appends the current date and time to the file name when it is uploaded to the FTP server.

The system allows you to specify two FTP server addresses. At upload time, the system tries to connect to the primary FTP server and, if it fails, it tries the secondary FTP server. You must configure the username and password for establishing the FTP connection.

Statistics collected

The following list details the type of statistics that are collected and reported for each class.

- Packet counters. Counters are maintained for the total packets forwarded, total bytes forwarded, total packets dropped, and total bytes dropped. These are counts accumulated since system boot-up or since the last clear.
- Detailed drop statistics. Counters are maintained for packets dropped due to insufficient buffers, packets dropped due to max queue size reached, packets dropped by policing and packets dropped by RED. This is useful to determine the cause of the packet drops.

The statistics that are uploaded are of two types:

- Interface statistics
- Class statistics

The interface statistics can be extracted from the interface parent classes (root-in and root-out). The interface statistics include the following:

- Interface type
- Interface name

- Sample Number
- Time
- Uptime
- Packets forwarded
- Packets received
- Bytes forwarded
- Bytes received
- Packets dropped (for in and outbound direction)
- Bytes dropped (for in and outbound direction)
- Packets dropped due to maximum queue size (for outbound direction)
- Packets dropped due to insufficient buffers (for outbound direction)
- Packets dropped due to RED (for outbound direction)
- Packets dropped due to policing (for inbound and outbound direction)

The class statistics that are uploaded are as follows:

- Class Name
- Sample Number
- Sample End Time
- Interface Name
- Parent Class
- Packets forwarded
- Bytes Forwarded
- Packet dropped in that class
- Bytes dropped
- Packets dropped due to maximum queue size
- Packets dropped due to insufficient buffers
- Packets dropped due to RED
- Packets dropped due to policing

All these statistics are added to the file in a tabular format. The interface statistics are appended first followed by the class statistics.

Auto QoS

Auto QoS simplifies the provisioning of QoS by applying a pre-configured QoS treatment, either for all interfaces in a system or for a particular interface, using a single global or per interface CLI command.

The Auto QoS configuration provides preferential treatment for high-priority traffic over regular or best effort traffic.

The following table describes the requirements for meeting the network service classes (NSC) implementation guidelines. When Auto QoS is enabled on chassis interfaces, the policy classes, queue priorities, queue shaping, and congestion parameters are provisioned to meet these NSC requirements.

The table describes how Auto QoS is going to behave for chassis interfaces.

Table 10: Network service classes

Network Control Traffic Category	Network Service Class	Target Applications	Service characteristics	
Network control	Critical	• Super user Telnet • Critical heartbeat between nodes	Loss tolerance	Very low
			Delay tolerance	Very low
			Jitter tolerance	N/A
			Traffic profile	Small variable-sized packets
			Diff-Serv Codes	CS7
			Queuing, Scheduling, and Shaping	CR: 10% PR: 100% CBQ Pri: 1 Priority Tail drop
	Network	• ICMP, OSPF, BGP, RIP, ISIS • SIP signaling between call servers in carrier networks • COPS, RSVP • DNS, DHCP, BootP, high priority OAM • Control and signaling between administrative domains	Loss tolerance	Low to very low
			Delay tolerance	Low
			Jitter tolerance	N/A
			Traffic profile	Variable-sized packets

Comments? infodev@avaya.com

Network Control Traffic Category	Network Service Class	Target Applications	Service characteristics	
			Diff-Serv Codes	CS6
			Queuing, Scheduling, and Shaping	CR: 10% PR: 100% CBQ Pri: 2 Weighted Tail drop
Interactive	Premium	• VoIP (G.711, G.729 and other codecs) • Telephony signaling between gateway or end device and call server (H.248, MGCP, H.323, SIP) • Lawful Intercept • T.38 Fax over IP • Voice-band data over IP (modem) • Circuit Emulation over IP	Loss tolerance	Very low
			Delay tolerance	Very low
			Jitter tolerance	Very low
			Traffic profile	Typically, fixed-sized packets
			Diff-Serv Codes	EF, CS5
			Queuing, Scheduling, and Shaping	CR: 35% PR: 100% CBQ Pri: 3 Priority Tail drop
	Platinum	• Interactive video (video conferencing) • Interactive gaming	Loss tolerance	Low
			Delay tolerance	Low
			Jitter tolerance	Low
			Traffic profile	Variable-sized packets
			Diff-Serv Codes	CS4, AF41, AF42, AF43
			Queuing, Scheduling, and Shaping	CR: 10% PR: 100% CBQ Pri: 4 Weighted CS4: 80-99% AF41: 80-99% AF42: 60-85% AF43: 50-80%
Responsive	Gold	• Streaming audio • Streaming video (video on demand) •	Loss tolerance	Very low to low

Network Control Traffic Category	Network Service Class	Target Applications	Service characteristics	
		Broadcast TV • Pay per view movies and events • Video surveillance and security • Webcasts	Delay tolerance	High
			Jitter tolerance	High
			Traffic profile	Variable-sized packets
			Diff-Serv Codes	CS3, AF31, AF32, AF33
			Queuing, Scheduling, and Shaping	CR: 10% PR: 100% CBQ Pri: 4 Weighted CS3: 80-99% AF31: 80-99% AF32: 70-90% AF33: 60-80%
	Silver	• Client/server applications • SNA terminal to host transactions (SNA over IP using DLSw) • Web-based ordering • Credit card transactions • Financial wire transfers • ERP applications (such as SAP/BaaN)	Loss tolerance	Low
			Delay tolerance	Low-medium
			Jitter tolerance	N/A
			Traffic profile	Variable-sized packets
			Diff-Serv Codes	CS2, AF21, AF22. AF23
			Queuing, Scheduling, and Shaping	CR: 5% PR:100% CBQ Pri: 5 Weighted CS2: 75-99% AF21: 75-99% AF22: 60-80% AF23: 40-70%
Timely	Bronze	• Store and forward applications • Email • Billing record transfer • Non critical OAM&P (SNMP, TFTP)	Loss tolerance	Low
			Delay tolerance	Medium
			Jitter tolerance	N/A

Comments? infodev@avaya.com

Network Control Traffic Category	Network Service Class	Target Applications	Service characteristics	
			Traffic profile	Variable-sized packets
			Diff-Serv Codes	CS1, AF11, AF12, AF13
			Queuing, Scheduling, and Shaping	CR: 10% PR: 100% CBQ Pri: 5 Weighted CS1: 75-99% AF11: 75-99% AF12: 60-80% AF13: 40-70%
	Standard	• All traffic not in any of the other classes • Best effort traffic	Loss tolerance	Typically not specified
			Delay tolerance	Typically not specified
			Jitter tolerance	N/A
			Traffic profile	Variable-sized packets
			Diff-Serv Codes	CS0, DF
			Queuing, Scheduling, and Shaping	CR: 9% PR:100% CBQ Pri: 8 Weighted 70-99%

To meet the NSC guidelines, two Auto QoS profiles are created: one for the inbound direction (AutoQoSPolicyIn) and the other for the outbound direction (AutoQoSPolicyOut). The appropriate QoS feature, either CBQ (shaping) or Policing, and direction is selected based on the interface type.

For WAN interfaces the CBQ feature is applied for the outbound direction.

The queuing, scheduling, and shaping parameters are tailored to use the existing CBQ functionality. The priority, bandwidth parameters, and RED thresholds are chosen to approximate the behavior in NSC and NSC-TE guidelines.

For Chassis Ethernet interfaces, the policing feature is applied for the inbound direction.

Important:

For PPP bundles, when auto QoS is enabled, if the bundle bandwidth is less than or equal to 768Kbps, the LFI feature is enabled automatically. The feature is disabled when auto QoS is removed.

Auto QoS class map for WAN

To identify various kinds of traffic, the NSC guidelines require eight class maps, which are identified using DiffServ code points. Auto QoS uses the hierarchical flow classification to define these eight QoS class maps on WAN interfaces.

The following diagram shows the AutoQoS hierarchical classification tree for most WAN interfaces.

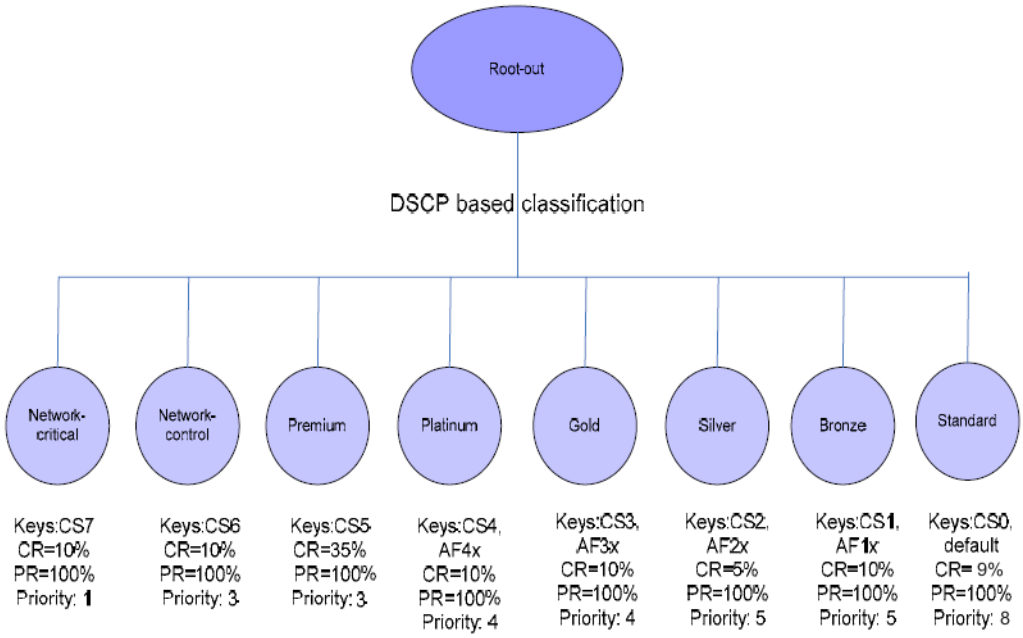

Figure 9: Auto QoS parameters for WAN interfaces

Auto QoS policy for Frame Relay and low bandwidth PPP interfaces

In addition to the two Auto QoS profiles, AutoQoSPolicyOut and AutoQoSPolicyIn, a special profile named AutoQoSPolicyOut4Q is also created during system initialization.

The primary reason for having this additional Auto QoS profile is to support low bandwidth bundle interfaces in the outbound direction. Since the AutoQoSPolicyOut profile has eight class maps in it, it requires a minimum of (8*4) 32 buffers plus 4 buffers for out-bound-default class (for a total of 36 buffers).

But if a bundle interface is configured with much less bandwidth (like 64kbps), it is assigned a minimum of only 20 buffers. With 20 buffers, the interface can service only 5 classes, so AutoQoSPolicyOut profile cannot be applied to these interfaces.

In these cases, when buffers are less than 36, AutoQoSPolicyOut4Q is applied for PPP interfaces.

And by default, for Frame Relay interfaces, AutoQoSPolicyOut4Q profile is applied in all cases (regardless of buffers).

Table 11: AutoQoSPolicyOut4Q

Class-name	DSCP values	CR	PR	Priority
network	cs7, cs6	20	100	1
voice	cs5, ef	35	100	2
oam	cs4, af41, af42, af43	10	100	4
best-effort	cs3,af31,af32,af33,cs2,af21,af22,af23,cs1,af11,af12,af13,cs0,default	34	100	4

Additional parameters for inbound Auto QoS

The following table shows the QoS parameters used for inbound Auto QoS.

Table 12: Inbound Auto QoS parameters

Class map	CBQ			Police	
	CR (%)	PR (%)	Priority	CIR	PIR
Critical	0	0	8	1%	-
Control	0	0	8	5%	-
Voice	0	0	8	50%	-
Video	0	0	8	10%	15%
Streaming	0	0	8	10%	-
Transaction	0	0	8	5%	-
Oam	0	0	8	15%	-
Best-effort	0	0	8	-	-

Enabling and disabling of Auto QoS

By default, auto QoS is disabled on all interfaces. To enable auto QoS on an interface, QoS must be enabled on the interface (either globally or at the interface level), but the interface

cannot have a manual QoS configuration applied. That is, the interface cannot be configured with a user-defined policy using the **service policy** command or have QoS features enabled using the **enable [monitoring | policing | cbq] <direction>** command. If the interface has any manual QoS parameters configured, then applying auto QoS on the interface fails.

You can enable the auto QoS feature by enabling auto QoS globally, or by enabling auto QoS at the interface level.

If you enable auto QoS at the global level, auto QoS is applied to all existing eligible interfaces, and to any interfaces that are subsequently created.

If you enable auto QoS at the interface level, the command applies to that interface, whether or not auto QoS is enabled or disabled globally.

You can disable the auto QoS feature by disabling auto QoS globally, or by disabling auto QoS at the interface level.

If, you disable auto QoS globally, auto QoS is disabled on all interfaces, whether or not auto QoS was enabled at the global or at the interface level.

Modifying Auto QoS policy maps

If you wish to adjust the default auto QoS settings to better suit your implementation, you can modify, but not delete, the existing auto QoS class maps. Any changes to the policy maps are propagated to all interfaces enabled with auto QoS.

Alternatively, you can clone the existing auto QoS policy maps, modify them as required, and then apply them to interfaces as normal.

Control Traffic prioritization

The Network Control (NC) traffic is not of interest to an end-user but is necessary for the network to operate properly. This network Control traffic is different from application control (signaling) traffic required for some user applications. Network switches or routers initiate the Network control traffic. For example, call setup SIP signaling traffic between an IP Deskphone and a call server that controls it, is considered as user traffic and not Network Control traffic. However, routing table updates being propagated across a network are considered Network Control traffic since they are not end-user applications.

Network Control traffic must use a queue that is separate from the user/data traffic if there are a sufficient number of queues available to support Network control traffic and the desired number of user services. Network Control traffic typically requires little bandwidth but must be assured that it gets transmitted across the network to keep the network operational. Network Control traffic must have a minimum guaranteed bandwidth.

Network administrators use Critical Control (NC) traffic to operate the network. User or application traffic must not be mapped into this queue which is dedicated for Network control traffic.

The rest of the sub-sections will deal with the details of constituents of Network Control traffic, how it gets prioritized, and its QoS treatment.

The various traffic which termed as network control traffic is used to keep the network operationally up. From QoS perspective we will categorize the network control traffic in three types. They are

- Layer 2 control traffic
- The PPP keepalive messages, which are necessary to keep PPP link operationally up. The ARP packets are considered as control traffic, if not otherwise results in congestion. Similarly VLAN bridge control traffic, like BPDU packets, to keep the bridge operationally up. In Frame Relay, LMI traffic considered as control traffic to keep logical PVC interfaces operationally up. In case of PPPoE, the discovery and PPP negotiation traffic are considered as a control traffic, which is required to establish and keep the PPP link operationally up.
- • Layer 3 control traffic
- The ICMP, DNS, and DHCP packets are considered as control traffic.
- Network control traffic
- The RSVP, BGP, RIP, OSPF, IGMP, IGRP, VRRP, PIM, and GRE keepalive messages are considered as network control traffic.

Chapter 5: Ethernet Module QoS fundamentals

The following figure shows the overview of features available with Ethernet module QoS flows.

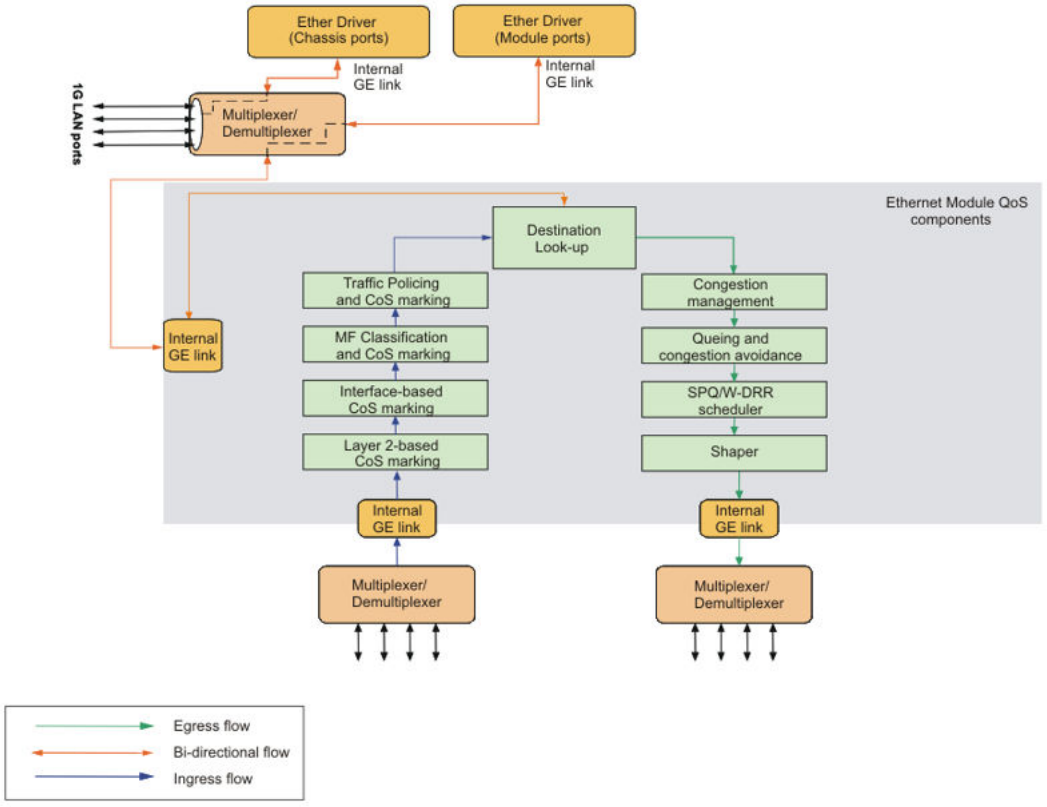

Figure 10: Overview of Ethernet module QoS operation

The following sections provide details for each of the Ethernet module QoS features.

Important:

On the Avaya Secure Router 4134, the GigE Large Ethernet module is oversubscribed and provides a 4:1 blocking ratio. That is, a bandwidth of one gigabit is shared by every four ports. Within a group of

four ports, any QoS treatment applied to one port is also applied to the three remaining ports in the group, with the exception of the following features, which are applied independently for each port:

- QoS status
- Auto QoS status
- Interface DSCP marking
- Interface UP marking
- Service policy attached to the interface (policy map and its associated configuration)

CoS marking

Every packet that enters the Ethernet modules is associated with a set of CoS attributes. The CoS attributes determine how the packet gets serviced and how the DSCP and 802.1p User-Priority values are set in the packet header on its exit from the Ethernet module. The Ethernet modules support the following four CoS attributes:

- User Priority (UP)

  The User Priority (UP) refers to the 802.1p user priority field in the VLAN or priority tagged packet header. The UP CoS marking is used to set the header field of the packet on its exit from the Ethernet module. The UP value ranges from 0 (lowest priority) through 7 (highest priority).

- Drop Precedence (DP)

  The Drop Precedence of a packet specifies the probability that the packet will be dropped during congestion. Packets with higher DP have higher probability of being dropped than those with a lower DP. The DP values can be low, medium, or high.

  The DP is used by congestion avoidance mechanisms at the egress QoS processing stage for making drop decisions. (See Congestion control and avoidance on page 76). In addition, the traffic policing engine can use this attribute as an initial conformance indicator with color-aware policing. (See Traffic policing on page 72).

- Egress queue

  There are 8 egress queues supported on each Ethernet Module port, each operating at different queue priority. The egress queue values range from 1 through 8, with 1 being the highest priority queue and 8 being the lowest priority queue. The egress queue marking determines the egress queue to which the packet is en-queued.

  For more detailed explanation of egress priority queues, refer to Queuing and Scheduling on page 77.

- DiffServ Code Point (DSCP)

  The DSCP refers to the DiffServ field (first 6 bits of the TOS/Traffic Class octet) in the IPv4/IPv6 packet header. The DSCP CoS marking is used to set the header field of the packet on its exit from the Ethernet module. This value determines the PHB (Per Hop Behavior) given to the packet at each downstream node in the DS domain.

The DSCP values range from 0 through 63. For a list of DSCP values and standard PHBs associated with them, see Table 4: IETF recommended DSCP values on page 27.

Egress queue and DP marking are internal CoS attributes used within the Ethernet module. UP and DSCP are external CoS attributes that are used within the Ethernet module and also used to mark the respective fields in the packet headers.

Ethernet module QoS performs CoS marking and re-marking at various stages during ingress QoS processing. Each processing stage may preserve or override the CoS attribute settings done by the pervious stages.

In the inbound direction, the following markings can take place, in the order listed:

1. Layer 2-based marking
2. Interface-based marking
3. Policy-based marking
4. Policing-based marking

The following sections describe how each of the four CoS attributes are marked or remarked at each stage of QoS processing.

Layer 2-based marking

With Layer 2-based marking, you can configure the system-wide mapping of Layer 2 UP values to DP attributes for tagged ingress packets. You can also configure the system-wide mapping of the UP values to specific egress queues.

The following table describes the default UP-based mappings for DP marking. When Ethernet module QoS is enabled, these mappings are applied by default. You can modify these system-level default mappings.

Table 13: New515

UP	DP
0, 1, 2, 3	high
4, 5, 6	medium
7	low

For tagged ingress packets, the egress queue is assigned based on the UP value in the Layer 2 header. The following table shows the default UP to egress queue mapping that is used for this purpose. You can modify these system-level default mappings.

Table 14: User priority to egress queue default mapping

UP	egress queue
0	8
1	7

UP	egress queue
2	6
3	5
4	4
5	3
6	2
7	2

UP marking based on egress queue value

As part of Layer 2-based CoS marking, Module QoS also provides a mechanism for marking the UP of untagged packets based on the egress queue.

For VLAN or priority tagged packets, the UP value attribute is copied from the 802.1p user-priority field. For untagged packets the UP is set based on egress queue.

The following table describes the egress queue to UP mappings. You can modify these default mappings.

Table 15: User Priority to Drop Precedence default mapping

egress queue	UP
8	0
7	1
6	2
5	3
4	4
3	5
2	6
1	7

Interface-based CoS marking and default egress queue assignment

For untagged ingress packets, you can configure per-port UP and DSCP marking. When you enable CoS marking on a port, all untagged ingress packets are remarked with the configured UP and DSCP attributes.

You can also assign a default egress queue on the port. Typically, the default queue is enabled for untagged packets, but if desired, you can force all packets, tagged and untagged, that enter the port to use the configured default queue.

Policy map-based marking

You can use policy-maps to configure the marking of CoS attributes based on the classified Layer2-4 header information. A policy map is a set of rules to classify traffic based on the Layer 2-4 packet header information. A set of rules can be associated with a class and there can be multiple rules within a policy map. Every class in the policy map has an associated set of actions. One available action is CoS marking. For any class, marking of any of the CoS attributes (DP, UP, DSCP, and egress queue) can be enabled or disabled.

For more details on policy maps, refer to Multifield traffic classification on page 67.

Policing-based CoS marking

Traffic policing determines the conformance level of a packet (conform/exceed/violate) based on its conformance to configured rate and burst size (see Traffic policing on page 72).

Policing based CoS remarking modifies the packet's CoS attributes based on its previously marked CoS attributes (DSCP for IP packets and UP for non-IP packets) and the assigned conformance level.

Multifield traffic classification

Similar to Chassis QoS, multifield traffic classification in Ethernet module QoS identifies and segregates traffic that matches particular criteria to classify aggregate traffic into separate flows or traffic classes.

Unlike Chassis QoS, the traffic classification on Ethernet module interfaces can be configured for inbound traffic only. However, the properties you define for an inbound traffic class can affect how the traffic is serviced in the outbound direction.

The available QoS options differ on inbound and outbound flows. Further, The QoS options and their implementation with Ethernet module QoS differ from those available with Chassis QoS.

Class map

To create class maps for multifield traffic classification, Ethernet module QoS uses a classification approach that differs from Chassis QoS classification.

Rather than allowing only one match rule per configured class as in Chassis QoS, each class in Ethernet module QoS can have multiple classification rules. The packet is considered to be successfully classified into a class when the classification process matches any one of the configured classification rules for that class.

Once you have specified the traffic classification rules, you can then configure desired QoS parameters to define the QoS treatment for the packets matching that class. The available QoS options are CoS marking and re-marking, policing, and policy-based redirect. (Other Ethernet module QoS features, including monitoring, congestion management, queuing and scheduling, and shaping, are enabled at the interface level, separately from the policy map configuration.) The actions associated with a class are applied for all packets that are classified into the class.

A policy map represents the complete set of classification rules that can be associated with one or more interfaces. The advantage of having multiple classes within a policy map is that you can associate different actions with different classes within the same policy map.

Class sequence

With Ethernet module QoS, the class map is not defined by a tree structure with parent and child classes. Instead, classification is list-based, with support for heterogeneous classification attributes.

Within each policy map, classes are listed in order of search priority. When searching for a match for a packet, Module QoS searches through the configured classes sequentially and matches the packet with the first matching class in the sequence. The action associated with this class is then carried out.

You can control the search sequence for classes by configuring the priority parameter for each class. The priority specifies the search sequence of each class from 1 to 1024.

If you specify the priority number of a new class to be the same as an existing class, the order of the class list dynamically changes so that the new class is inserted at the specified priority, and the classes of equal and lower priority move down one level to accommodate the new class.

If you do not specify a priority number for a class, it is added to the end of the list of classes in the policy map.

Rule search sequence

Within each class, the match rules are also listed in order. The rules are searched sequentially from start to end and the matching process terminates when a matching rule is found. For this reason, it is advisable to keep more specific rules higher in the order than more generic rules.

Once a matching rule is found, the associated actions configured for the class are carried out.

Unlike classes, no parameter is available to specify the priority order of match rules within a class. Instead, the rules are searched in the same sequence in which you enter them. The first rule you enter for a class becomes the first rule in the list to be searched. As a result, the order in which you create the rules in a class is important.

Sample class map

The following figure show a sample class map to illustrate how classes and rules are ordered and searched within a policy map.

```
class-map: critical (seq# 1)
   0 bytes forwarded
   queue 1, drop-precedence carry, dscp carry, user-priority carry
   rate-monitoring:
      disabled
   rules:
      match ipv4 dscp cs7
      match ipv6 dscp cs7
   police:
      disabled
   redirection:
      none

class-map: control (seq# 2)
   0 bytes forwarded
   queue 2, drop-precedence carry, dscp carry, user-priority carry
   rate-monitoring:
      disabled
   rules:
      match ipv4 dscp cs6
      match ipv6 dscp cs6
   police:
      disabled
   redirection:
      none
```

Figure 11: Sample class maps

If this policy map is applied on an interface, and then a packet with, for example, an IPv6 DSCP value of CS6 arrives on the interface, Ethernet module QoS searches through the class maps sequentially to find a match. Starting with class map 1, the search moves from the first rule "match ipv4 dscp cs7" to the second rule "match ipv6 dscp cs7". Finding no matches, the search moves to class map 2, starting with "match ipv4 dscp cs6". When a match is found at the second rule, "match ipv6 dscp cs6", the action for this class map is carried out.

Policy map and QoS actions

For each class that you define, you can associate QoS parameters to that class. The possible QoS features that you can apply to the classes are as follows:

Inbound only

- CoS marking
- policing (srTCM/trTCM)
- policy-based redirection

As Ethernet module QoS supports policy maps on the inbound direction only, you do not need to specify a direction for the policy map.

The Ethernet module QoS supports additional QoS features in the outbound direction, namely SP/WRR (queuing, scheduling), shaping, and congestion management. However, these

features are configured and enabled directly at the port and queue level rather than with policy maps.

Classification attributes

The MF classification of incoming packets is supported based on any combination of the following Ethernet, IPv4 and IPv6 packet fields:

- Ethernet MAC Source Address
- Ethernet MAC Destination Address
- Ethernet Type
- VLAN ID
- 802.1p value (User Priority)
- IPv4/IPv6 Source address (or prefix)
- IPv4/IPv6 Destination address (or prefix)
- IPv4 Protocol Type or IPv6 Next Header
- IPv4/IPv6 Source/destination TCP/UDP ports
- IPv4/IPv6 DSCP
- IPv4 Precedence
- IPv4 ToS
- IPv6 Traffic Class
- IPv6 Flow Label

Classification field values in a matching rule can be specified either as single value or as a range of values depending on the kind of field used for the classification. The following table describes the possible match operations for each of the fields in a matching rule:

Table 16: Classification match operations

Field	Match Operation
Ethernet MAC Source Address	EQUAL, RANGE, NEQ
Ethernet MAC Destination Address	EQUAL, RANGE, NEQ
Ethernet Type	EQUAL, NEQ
VLAN ID	EQUAL, RANGE, NEQ
802.1p value (User Priority)	EQUAL, NEQ
IPv4 Source address	EQUAL, RANGE, NEQ, MASK
IPv6 Source address	EQUAL, RANGE, MASK

Field	Match Operation
IPv4 Destination address	EQUAL, RANGE, MASK, NEQ
IPv6 Destination address	EQUAL, RANGE, MASK
IPv4 Protocol Type	EQUAL, RANGE, NEQ
IPv6 Next Header	EQUAL, RANGE, NEQ
IPv4 Source ports	EQUAL, RANGE, NEQ, GTE, LTE
IPv4 Destination ports	EQUAL, RANGE, NEQ, GTE, LTE
IPv4/IPv6 DSCP	EQUAL, RANGE, NEQ
IPv4 Precedence	EQUAL, RANGE, NEQ
IPv4 TOS	EQUAL, RANGE, NEQ
IPv6 Traffic Class	EQUAL, RANGE, NEQ
IPv6 Flow Label	EQUAL, RANGE, NEQ

Classification types

When you create rules within a class map, you must specify the rule type for each rule. The following three types of classification rules are supported based on the packet type.

- Non-IP rules – defined based on Ethernet fields only
- IPv4 rules – defined based on IPv4 and Ethernet fields
- IPv6 rules – defined based on IPv6 and Ethernet fields.

Assigning a policy map to interfaces

After you have configured a policy map, you can associate it with one or more interfaces. The policy map can be associated with an interface in the inbound direction only. When you associate a policy map with an interface, an instance of the policy map is instantiated over the interface in the specified direction.

After you associate a policy map with an interface, unlike Chassis QoS, you do not need to enable the interface with the QoS features that are specified in the policy map. The specified features are automatically enabled.

Modifying a policy map

If you make any changes to a policy map that is already associated with one or more interfaces, then all applied changes are propagated to all the previously associated interfaces.

Deleting a policy map

If a policy map is currently associated with an interface, you cannot delete that policy map. Before you can delete the policy map, you must first remove it from all applicable interfaces.

Policy map clones

You can clone a policy map and save it to another policy map name to create a policy map that has minor differences from an existing configuration.

Traffic policing

Traffic policing is one of the available actions that you can apply to configured traffic classes.

Traffic policing meters traffic and performs actions based on the results of the metering. Metering compares the properties of the classified flows against a traffic profile and categorizes the traffic flows accordingly as conformed, exceeded, or violated. The traffic is then tagged as green, yellow, or red respectively.

With Ethernet module QoS, policing actions are based not only on the conformance level of the packet, but also on the incoming DSCP value (for IP packets) or UP value (for non-IP packets) of the packet. Packets that exceed the profile can be remarked (DSCP or UP), and the packets that violate the profile can be dropped.

Policing can be enabled on the ingress of any Ethernet Module interface. You can apply policing to any classes that are identified by the Multifield traffic classifier.

Ethernet module QoS supports the same metering algorithms as in Chassis QoS:

- srTCM
- trTCM

Both policing methods use a token bucket algorithm, where a token is considered as a byte and the bucket is a counter that is updated during the token refilling mechanism.

Policing using Single Rate Three Color Marker

srTCM uses a single rate Committed Information Rate (CIR) and two burst sizes, Committed Burst Size (CBS) and Excess Burst Size (EBS). The srTCM is implemented using two token buckets, namely Tc of size CBS and Te of size EBS. Both the token buckets are filled at the same rate (CIR).

When you configure srTCM, you must specify the CIR value. The CIR determines the token fill rate. You can also specify CBS and EBS but they are not mandatory values. By default CBS is one second worth of data corresponding to CIR and EBS is twice CBS. For example if CIR is configured as 10Mbps then CBS is 10 Mb and EBS is 20 Mb.

A color of green is assigned by the meter if the packet does not exceed the CBS, yellow if it does exceed the CBS but not the EBS, and red if it exceeds the EBS. Coloring is only a means to convey the conformance level of packets. An action can be configured for each color. The supported actions include: Mark DSCP value, Mark UP value, and Drop.

srTCM operation

For a detailed description of srTCM, refer to Policing using Single Rate Three Color Marker on page 29. Be aware that, in addition to supporting DCSP remarking similar to Chassis QoS, Ethernet Module QoS also supports remarking of UP.

Policing using Two Rate Three Color Marker

The trTCM feature also utilizes two token buckets namely Tc of size Committed Burst Size (CBS) and Tp of size Peak Burst Size (PBS). But in this case, each token bucket has a different token fill rate, Committed Information Rate (CIR) and Peak Information Rate (PIR) respectively. You can configure the CIR, PIR, CBS and PBS.

A color of green is assigned by the meter if the packet doesn't exceed the CBS, yellow if it does exceed the CBS but is less than or equal to the PBS, and red if it exceeds the PBS. A policing action can be configured for each resulting color. Based on the conformance level, the associated action is performed on the packets. Supported actions include Permit, Mark-DSCP and Drop.

When you configure trTCM, you must specify the CIR and PIR values. You can also specify CBS and PBS but they are not mandatory values. By default, CBS is one second worth of data corresponding to CIR and PBS is one-second worth of data corresponding to PIR. For example, if CIR is configured as 10Mbps and PIR is configured as 20Mbps, then CBS is 10 Mb and EBS is 20 Mb.

trTCM operation

For a detailed description of trTCM, refer to Policing using Two Rate Three Color Marker on page 31. Be aware that, in addition to supporting DCSP remarking similar to Chassis QoS, Ethernet Module QoS also supports remarking of UP.

Policing actions

For each metering color, you can specify the policing actions. For conforming (green), exceeding (yellow), and violating (red) packets, you can configure remarking of CoS attributes

(DSCP and UP). However, unlike Chassis QoS, you can only configure dropping of packets for violating (red) packets. All other packets are allowed.

Policing based CoS remarking values cannot be configured for each class. Only system level configuration is supported.

Color aware mode for srTCM and trTCM

Color aware mode is useful if the classified packet is already conditioned with color information. Unlike Chassis QoS, where DSCP information is used to obtain the incoming color information, Ethernet module QoS uses the DP value.

In this case, a DP value of low corresponds to conformance color green, medium corresponds to conformance color yellow and high corresponds to conformance color red.

When srTCM or trTCM are enabled with color aware mode, they consider each packet as pre-colored packet before policing it as normal.

The color aware mode can be enabled or disabled on per-class basis.

Accounting

The Ethernet module supports flow based accounting which can be enabled or disabled for every class. Accounting can be enabled or disabled at the system level also. The following accounting statistics are supported:

- Number of octets forwarded
- Number of conforming octets
- Number of exceeding octets
- Number of violating octets

Policy-based redirect

Policy-based redirect is one of the available actions that you can apply to configured traffic classes.

With policy-based redirect, you can specify that packets from a specified class are to be forwarded to a destination specified by any one of the following parameters:

- IPv4 address of the next hop
- IPv6 address of the next hop

- Interface name (bundle name or tunnel name)
- MPLS LSP name

Traffic monitoring

Unlike Chassis QoS, on Ethernet modules, there is no restriction on enabling traffic monitoring with other QoS features such as policing and shaping.

The network processor supports flow based packet statistics collection. This is done using the billing counters associated with a policer that can be associated with a flow. The statistics available are

- Number of forwarded octets
- Number of forwarded green octets
- Number of forwarded yellow octets
- Number of forwarded/dropped red octets

These counters can also be used to measure the rate of any flow. The input rate of a flow or a flow aggregate can be measured by sampling the sum of green, yellow and red octet counters and the policer output rate can be measured by sampling the forwarded octets counter. These statistics serve as very useful tools for you to study the behavior of traffic flows/flow aggregates over any given period of time.

The monitoring feature allows you to measure the rate of any flow or flow aggregate associated with a class. To measure the rate of a flow, you must do the following :

- Enable rate monitoring on a class associated with the flows to be monitored
- Define the sampling interval and the sampling period at the global level.
- Start the sampling process for all the classes using a single command.

The average flow rate calculation is done based on the moving average of the flow rate at every sampling interval. The commonly used method is based on the Exponentially Weighted Moving Average (EWMA) calculations. The following formula is used for calculating the average flow rate

FlowRateavg (t) = FlowRateavg (t-1) + (FlowRatecurr(t) – FlowRateavg(t-1)) * Wq

Where FlowRatecurr(t) is the flow rate at a specific sampling interval, FlowRateavg (t) is the average rate calculated in that sampling interval and FlowRateavg(t-1) is the average rate calculated in the previous sampling interval.

Congestion control and avoidance

The purpose of congestion avoidance is to identify the onset of congestion and start dropping packets quickly enough to avoid complete depletion of resources. Ethernet module QoS supports two packet dropping mechanisms, namely tail drop and Weighted Random Early Discard (WRED) with three levels of drop precedence (DP).

The tail drop mechanism simply drops packets when the number of packets queued exceeds a configured maximum limit.

With RED, the basic operating philosophy is to detect the onset of congestion and start dropping packets in a random fashion before queue overflow leads to tail drops. The result of the random drop is that the sending endpoint detects the dropped traffic and slows the TCP sending rate by adjusting the window size. Random drops not only improve throughput of adaptive applications using TCP but are also more suitable than tail drops for voice and streaming audio or video, as they result in less perceptible degradation.

RED makes decisions based on average queue length rather than instantaneous queue length. The Exponential Weighing Factor (EWF) is a user-configurable parameter that allows you to control the average queue size.

The WRED algorithm is an extension to RED in that it also takes into consideration the conformance level (DP) of the packet, so that traffic with high DP is dropped first. WRED supports three levels of drop precedence (DP) for handling congestion on egress queues: low, medium, and high (equivalent to green, yellow, and red respectively on Chassis QoS WRED). When congestion occurs on an egress queue, packets with higher DP can be configured to drop, while packets with low DP are queued.

In the case of WRED, all the congestion avoidance parameters (minth, maxth and mpd) are used to derive the drop curve while in the case of Tail Drop, only the maxth parameter is used to make dropping decisions.

Ethernet module QoS does not support DS-RED.

Unlike Chassis QoS, you cannot configure the congestion parameters separately for each queue. Instead, Ethernet module QoS supports four congestion configuration profiles. Each profile has congestion avoidance parameters per queue and per DP. For each drop precedence (DP) and queue, a set of congestion avoidance parameters comprising of maximum threshold (maxth), minimum threshold (minth), maximum probability denominator (mpd) are configured. An exponential weighted moving factor (EWF) is also configured per queue.

You can assign one of the four configured congestion profiles to any of the egress ports.

For more details on WRED, refer to Congestion control and avoidance with RED on page 35

Queuing and Scheduling

Ethernet modules support eight egress queues per port, and each queue is numbered from 1 through 8. Queue 1 has the highest priority and queue 8 has the lowest priority.

You can schedule the servicing of queues on each port using one of the following methods:

- Deficit Weighted Round Robin (DWRR)

  DWRR scheduling supports a guaranteed minimal bandwidth, as required for DiffServ Assured Forwarding (AF) support, while reducing the burstiness and jitter that is associated with non-WRR scheduling.

- Strict Priority (SP)

  Strict priority scheduling supports minimal latency for real-time traffic and provides DiffServ Expedited Forwarding (EF) support.

- A combination of SP and D-WRR

  With a combination of SP and DWRR scheduling, SP queues provide minimal latency for voice and mission critical protocols. The DWRR queues can then share the remaining link bandwidth, according to their relative weights.

DWRR

When you configure a queue to operate in DWRR mode, the queue is placed in a DWRR scheduling group.

Within the DWRR group, the queues share the bandwidth and are serviced according to their configured queue weights. Packets in a queue are scheduled for transmission unless a higher priority queue within the WRR group has traffic and has not exhausted its weight in one round of servicing.

Queues contained within a DWRR group must be continuous. For example, you cannot configure queues 6 and 8 in the same DWRR group without also including queue 7.

You can configure up to two separate DWRR groups.

Strict priority

Strict Priority queuing operates in an interrupt fashion. Traffic in higher priority queues is always scheduled prior to traffic in lower priority queues. If the highest priority queue contains frames, all processing in the lower priority queues is stopped, and the switch transmits the highest priority frames until that queue is empty.

When you configure a queue to operate in SPQ mode, queues are scheduled according to the queue number, starting with the highest queue 1, with decreasing priority through queue 8.

When the highest priority queue is empty, frames from the next priority queue, if any, are sent, in succession, from the highest priority queue to the lowest priority queue.

You can configure any queues with SP queuing, and the SP queues can be non-continuous.

Strict priority and DWRR

You can configure the egress queues such that some queues operate in SPQ mode and others are part of one or two WRR groups.

With the combination of strict priority and DWRR, all queues are processed in SP fashion according to their priority from 1 to 8, except when they belong to a DWRR group. In this case, all SP queues that are higher in priority are serviced before any queues in the DWRR group. When no higher priority queues have packets to send, then the queues in the DWRR group share the available bandwidth according to their configured weights.

Similarly, if the queues in a DWRR group are in a higher priority than a lower priority strict queue, all DWRR queues must be serviced before the lower priority strict queue can send packets.

For example, in the following figure, queues 3-5 have a higher priority, and are therefore serviced before queue 6, even though queue 6 is an SP queue.

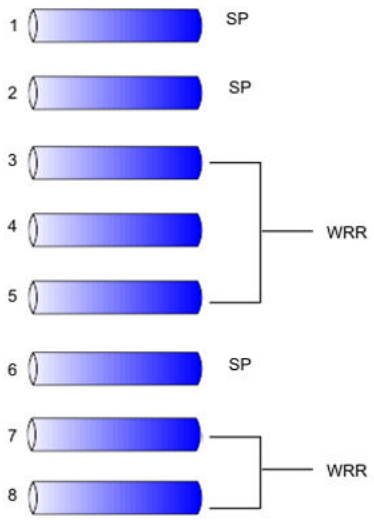

Figure 12: Strict priority and DWRR groups

Default queue configuration

By default, the scheduling mode of queues is a combination of both SPQ and WRR groups. The default configuration is as follows:

- Queues numbered 1, 2, 3, and 4 are in SPQ mode.
- Queues 5 and 6 belong to WRR group 1, and queues 7 and 8 belong to WRR group 0.
- The WRR group 0 queue weights are 2 and 1 for queues 8 and 7, respectively. The WRR group 1 queue weights are 2 and 1 for queues 6 and 5, respectively. Accordingly, these queues share according to their weights the remaining bandwidth not used by SP queues.

You can modify the default scheduling method and the parameters configured for each of the queues in Ethernet Module ports.

Traffic shaping

Traffic shaping smoothes out traffic bursts by queuing packets for a required amount of time. It limits a flow to a desired rate and burst. Ethernet module QoS supports shaping using a simple token bucket mechanism characterized by a peak rate and a burst size.

Ethernet module QoS supports shaping both at the port level and at the queue level. Shaping can be enabled independently for ports and queues. The configurable token bucket size determines the maximum burst size that can be permitted. The range of values is 4KB through 16MB in increments of 4KB. Each token corresponds to one byte. The shaping rates range between 651 Kbps and 99936 Kbps in multiples of 651 Kbps.

The token bucket is filled at the given peak rate, up to the configured maximum burst size. Packets are transmitted only if enough tokens are available in the token bucket. Tokens are depleted when packets are transmitted.

One of the inherent problems with shaping is the introduction of delay. Hence shaping is not suitable for real time applications like voice.

Buffer Management

The Ethernet modules utilize a store-and-forward architecture. Each packet received from a local network port or from the CPU is allocated a packet buffer from the local DRAM. Each packet buffer is associated with one or more packet descriptors (PD). A PD is allocated for each packet egress destination. A packet queued to be sent out on N egress ports has N

descriptors, all pointing to the same buffer in memory. Only after all the PDs are released will the packet buffer be released.

The Ethernet modules provide a total of 4000 packet buffers in each device. Each Ethernet module supports per port buffer limits in the ingress direction (receive buffers) and per-port packet descriptor limits in the egress direction to avoid over utilization of buffers by any specific port. The Ethernet module also supports packet descriptor limits at the queue level to allow fair sharing of buffers and packet descriptors across the queues associated with a port.

The network processor also supports flow control using 802.3x standards. The flow control mechanism is based on the current buffer occupancy levels at any given point in time.

The XOFF limit defines the threshold for generating 802.3x XOFF packets on an interface. If the total number of allocated buffers in the device exceeds this value, 802.3x XOFF packets will be sent out on the ingress interface at regular intervals. The XOFF packets force the preceding nodes to hold transmission for a specific amount of time.

The XON limit defines the threshold for generating 802.3x XON packets on an interface. If the total number of allocated buffers in the device falls below this value, 802.3x XON packets will be sent out on the interface. The XON packets trigger the preceding nodes to start any transmission that was put on hold because of previous XOFF indications.

Auto QoS

Please refer toAuto QoS on page 54 for the overview of the NSC guidelines that are used with Auto QoS feature. Auto QoS with the Ethernet module follows the same guidelines.

Auto QoS allows you to provision interfaces with default policies to support QoS for well known service classes. When Auto QoS is enabled on an interface, the default policy is mapped to that interface with rules to classify packets based on pre-marked DSCPs. The policer configuration is decided based on the deployment. The policy is configured to direct flows to the right queues based on their service classes.

Enabling and disabling of Auto QoS

By default, auto QoS is disabled on all interfaces. To enable auto QoS on an interface, QoS must be enabled on the interface (either globally or at the interface level), but the interface cannot have a policy map previously applied. That is, the interface cannot be configured with a user-defined policy using the **`service-policy`** command. If the interface has an existing policy map applied, then applying auto QoS on the interface fails. Other than applied policy maps, auto QoS is independent of the other Ethernet Module QoS features.

You can enable the auto QoS feature by enabling auto QoS globally, or by enabling auto QoS at the interface level.

If you enable auto QoS at the global level, auto QoS is applied to all existing eligible interfaces, and to any interfaces that are subsequently created.

If you enable auto QoS at the interface level, the command applies to that interface, whether or not auto QoS is enabled or disabled globally.

You can disable the auto QoS feature by disabling auto QoS globally, or by disabling auto QoS at the interface level.

If, you disable auto QoS globally, auto QoS is disabled on all interfaces, whether or not auto QoS was enabled at the global or at the interface level.

Modifying Auto QoS policy maps

If you wish to adjust the default auto QoS settings to better suit your implementation, you can modify the existing auto QoS class maps. However, you cannot delete any of the auto QoS class maps. Any changes to the policy maps are propagated to all interfaces enabled with auto QoS.

Alternatively, you can clone the existing auto QoS policy maps, modify them as required, and then apply them to interfaces as normal.

Enabling and disabling QoS

The QoS related processing for a packet is done at multiple stages on the NP packet processor. One of the important stages is the policy application stage (involving classification, policing and marking). The operator has the flexibility to define policies through the management interface (CLI). The operator can choose to bypass the policy application stage completely, if required. The QoS enforcement stages (involving congestion management, scheduling, and so on) can also be managed by the operator. However every packet transiting the NP will get subjected to QoS enforcement. It cannot be bypassed.

Using the QoS status management support, the operator can

- enable / disable policy application on any given interface
- disable policy application globally
- reset all the QoS parameters to default

Details

When the operator disables QoS globally, the following actions will be taken

- All interface-policy mappings will be removed from the Module
- If auto QoS is enabled on an interface, it would be disabled and the associated policy mapping will be removed
- All the QoS related parameters will be reset to default in the Module

When the operator enables QoS globally, all the QoS related parameters are re-instated in the Module.

When the operator disables QoS on an interface, the following actions will be taken

- The interface-policy mapping will be removed from the Module
- If auto QoS is enabled on that interface, it will be disabled and the associated policy mapping will be removed.
- Ingress marking parameters will be reset to default
- Shaping and scheduling parameters will be reset to default if the interface is not an over-subscribed interface. For over-subscribed interfaces, this will not be done since other interfaces sharing the cascade port will be affected

When the operator enables QoS on an interface, all the QoS related parameters for that interface are re-instated in the Module.

Disabling QoS on the Module only resets the QoS related parameters to default and does not switch off QoS processing entirely.

Ethernet Module QoS on the Avaya Secure Router 2330

On the SR2330, Ethernet module QoS refers to the QoS on all front-panel Ethernet ports (FE and GE).

The operation of Ethernet module QoS on the SR2330 is similar to the operation on SR4134 Ethernet modules. However, there are certain differences on the SR2330 hardware. The following sections describe these differences.

Class map restrictions

The SR 2330 chassis imposes additional restrictions on configuration combinations of rule types in a given class map. Specifically, no class map can have mixed rules of non-IP, IPv4, and IPv6.

The SR 2330 supports the configuration of exact match rules only. Range, negation and operators are not supported.

Global accounting enable and disable not supported

The SR2330 does not support global enable and disable of the accounting feature. Accounting and billing can be enabled and disabled for each class map, similar to the SR4134.

```
SR2330/configure/qos/module# ?
  accounting      Enable or disable accounting on LAN
                  modules. This command is not supported.
```

policing-cos-map and queue-cos-map not supported

The SR2330 does not support the global configuration tables **policing-cos-map** and **queue-cos-map**.

```
SR2330/configure/qos/module# ?
  policing-cos-   Configure policing based CoS remarking.
  map             This command is not supported
  queue-cos-map   Configure default output queue based CoS
                  marking. This command is not supported
```

To configure policing CoS marking, use the **remark-cos** command in the policing context (policy-map configuration), as shown in the following example.

```
SR2330/configure/qos/module/policy-map policy1/class-map
class1/police# remark-cos conform dscp af10 drop-precedence
low user-priority 6
```

Queue mapping is derived from the **user-priority-cos-map** table using the UP value obtained from either the incoming packet or from the interface default UP configuration. The **user-priority-cos-map** table, which you can configure using the **user-priority-cos-map** command under **qos/module**, is shown in the following sample output.

```
SR2330/configure/qos/module# show qos module system
qos enabled, auto-qos disabled, accounting enabled rate-
monitoring stopped, 60 sec sampling-interval, 600 sec
sampling-period
user-priority-cos-map:
user       output      drop precedence
priority   queue
0          8           high
1          7           high
2          6           high
```

```
3           5           high
4           4           medium
5           3           medium
6           2           medium
7           2           low
```

WRED not supported

The SR2330 only supports two congestion management schemes: Simple Random Early Drop (SRED) and Tail drop. Unlike the SR4134, WRED is not supported. As a result, the command **exponential-weighting-constant**under **congestion-profile** configuration is not applicable and not supported.

```
SR2330/configure/qos/module/congestion-profile 1# ?
  exponential-      Configure Weight for calculating moving
  weighting-        average of queue occupancy. This command is
  constant          not supported
```

Congestion limits for red and yellow packets only

Ethernet Module QoS in SR2330 only supports congestion limits for RED and YELLOW colored packets. For GREEN colored packets, the threshold is always the maximum descriptor limit. As a result, the values entered for **low** Drop Precedence (that is, GREEN colored packets) in the **drop-curve** command are ignored.

```
SR2330/configure/qos/module/congestion-profile 1# drop-curve
1?
  low               Low drop precedence. This parameter is
                    ignored.
```

Port-level queue assignment not supported

The SR2330 does not support port-level queue assignment for tagged packets, so the **default-queue** command is not supported. Output queue for packets ingressing through the interface is obtained from the **user-priority-cos-map** table using the UP value

obtained from either the incoming packet, if tagged, or from the interface default UP configuration (configured using the **mark-user-priority** command), if untagged.

```
SR2330/configure/interface/ethernet (0/4)/qos/module# ?

  default-queue    Assign default output queue for all
                   incoming packets on this interface. This
                   command is not supported. Use mark-user-
                   priority command instead.

  mark-user-       Mark user priority for incoming un-tagged
  priority         packets. This determines output queue for
                   the packets
```

Excess buffer limit configuration not supported

The SR2330 does not support excess buffer limit configuration for an interface. The packet processor provides 1024 egress buffer descriptors dedicated to each FE port and 1536 descriptors dedicated to each GE port. These descriptors can be distributed among the 8 queues for each port. This distribution of transmit buffers and descriptors among the interface queues is configurable using the **queue-limit** command, as shown in the following example.

```
SR2330/configure/interface/ethernet (0/4)/qos/module/queue
1# queue-limit 90
```

Only one WRR group for interface queues

The SR2330 supports only one Weighted Round Robin (WRR) group for the interface queues unlike the SR4134 which supports two WRR groups. As a result, on the SR2330 the **wrr-group** option in the **wrr-queue** command is ignored.

```
SR2330/configure/interface/ethernet (0/4)/qos/module/queue
1# wrr-queue 24 ?

  0 - 1            WRR group. This parameter is ignored.
```

Default marking of DSCP values at interface level not supported

The SR2330 does not support default marking of DSCP values at the interface level, so the **mark-dscp** command is not supported.

```
SR2330/configure/interface/ethernet (0/4)/qos/module# ?
  mark-dscp       Mark DSCP for all incoming packets. This
                  command is not supported
```

The global **dscp-cos-map** mapping table is used instead. You can view this table using the **show qos module dscp-cos-map** command.

Chapter 6: SLA fundamentals

SLA is a performance monitoring and measuring tool that you can use to monitor and measure network service performance between two nodes.

SLA operates system wide and is independent of interfaces. When you set the SLA destination properties, SLA uses any available interface to communicate with the destination, whether Chassis interface or Ethernet module interface.

SLA can operate continuously, periodically, or at a scheduled time.

The SLA functionality allows you to obtain the following performance measurements:

- Response Time or Round Trip Time
- Packet one-way delay
- Packet Jitter
- Data/Packet loss

The following diagram provides an overview of the SLA operation.

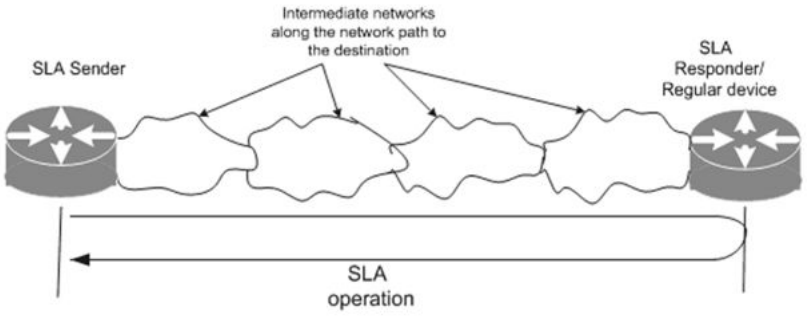

Figure 13: SLA operations

The main components of the SLA operation are the SLA Sender and the SLA Responder. The SLA sender resides in the network node where SLA performance attributes are measured. The SLA responder resides on a destination node separated from the SLA sender by one or more networks. Using SLA, you can measure the performance of the network or networks separating the SLA sender and responder.

The SLA sender functionality resembles that of a client in a client-server communication model. As with any client-server topology, the SLA sender sends requests to the SLA responder (the server) expecting a reply.

The SLA responder functionality resembles that of a server in a client-server model. The SLA responder listens on known dynamic and private port 50001 (50002 for IPv6).

When the SLA responder receives SLA data requests, it converts the request into a response and puts the appropriate information in the SLA data portion, based on the SLA type and operation fields in the

SLA header. The SLA responder also provides detailed timestamp information in the response. It then forwards the response to the sender.

SLA operations

SLA runs at the application level, using socket interfaces for transmitting and reception of SLA data packets, and UDP as a transport mechanism. SLA is independent of interfaces and operates system wide.

SLA can perform the following basic operations:

- UDP Echo
- UDP Jitter
- UDPv6 Echo
- UDPv6 Jitter
- ICMP Echo
- ICMPv6 Echo

The total number of SLA operations configured for the same destination is restricted to 5. The 5 entries can be of any SLA type for the same destination (even same types can be repeated). This restriction is applicable for both IPv4 and IPv6 destinations.

UDP and UDPv6 Echo operation

This operation is used to measure the end-to-end response time between the SLA sender and a specified destination location. The end-to-end response time is measured as the time difference between sending the UDP echo request and receiving the UDP echo reply.

With an Avaya Secure Router 2330/4134 SLA responder at the destination, the destination processing time is taken into account by the echo operation. In this case, the CPU time required to process the packet and the queuing delay are both removed from the response time.

If there is no SR2330/4134 SLA responder at the destination, the processing time cannot be removed from the response time calculation.

To initiate the UDP echo operation, you must provide the destination IP address, and optionally the private port number to which the SLA responder is listening.

By default, the SLA sender and responder both use the local port 50001 for IPv4 and port 50002 for IPv6.

Comments? infodev@avaya.com

You can also specify the port as port 7, which is the echo-server port. If the SLA responder is not an SR2330/4134 SLA responder, then you must specify port 7.

UDP and UDPv6 Jitter operation

The UDP jitter operation is used to measure performance sensitive metrics, including:

- jitter (variance of inter-packet delay)
- one-way delay (clock synchronization required)
- packet loss

Jitter computation

The packet jitter is defined as the variance in inter-packet delay.

To calculate the jitter, the SLA sender sends a series of timestamped sequenced packets to the destination at ten millisecond intervals (or at userconfigured interval). At the destination node, the SLA responder processes these packets, inserts the timestamps, and then sends the packets back to the source SLA.

Ideally, the destination receives the packets at the same interval as they were transmitted. However, delays in the network (such as queuing) can cause inter-packet arrival delay of greater or less than the configured interval. With positive jitter, packets arrive at longer intervals. Negative jitter occurs when the interval is shorter than originally configured. For example, if the configured interval is ten milliseconds, and the packets arrive 12 milliseconds apart, then the positive jitter is equivalent to two milliseconds.

The sender uses the timestamp information from the received packets to calculate the jitter, and one-way delay. The packet sequencing allows for calculation of the packet loss. (The packet loss is direction insensitive: if a packet is lost in either direction, then it will be accounted for as packet loss at the source.)

The jitter is calculated based on the algorithm given in RFC 1889 for packet interarrival jittery calculation.

To properly caclulate the jitter and one-way delay, the SLA sender and responder must have their clocks synchronized. SR2330/4134 systems use SNTP for clock synchronization.

ICMP and ICMPv6 Echo

This operation is used to measure the end-to-end response time, or connectivity to the specified destination. The response time is measured as the time difference between sending the ICMP or ICMPv6 echo request and receiving the echo reply. Since they use standard ICMP and ICMPv6 request and reply packets, the ICMP and ICMPv6 echo operations function similarly for all kind of responders (even non-SR2330/4134 systems).

SLA profiles

To configure an SLA operation, you must first configure an SLA profile. Each SLA profile can hold a single operation.

SLA supports a maximum of 1000 configured SLA sets or profiles.

Once you have associated an SLA profile with an operation, you can schedule the profile to run at specific times or intervals.

Actions and thresholds

With SLA, you can determine whether specific SLA thresholds related to jitter, delay, response time, and packet delay, among others, are being met or exceeded.

Chapter 7: Chassis QoS configuration

The high-level configuration steps for Chassis QoS are:

1. Create the policy map.
2. Configure the class-maps to classify the traffic.
3. Associate actions with the configured classes.
4. Apply the policy map to the interface in the appropriate direction.
5. Enable the specified QoS features on the interface.

Configuring multifield traffic classification

Creating policy map

Create a policy map to allow for the creation of class maps and associated actions.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To create the policy map, enter:

   ```
   [no] policy-map <policy-name> [force]
   ```

Table 17: Variable definitions

Variable	Value
<policy-name>	Specifies the name of the policy map.
[force]	When specified with the **no** parameter, forces the removal of the specified policy map from all applicable interfaces before deleting the policy map. If you do not specify this parameter, you must first manually remove the policy map

Variable	Value
	from all applicable interfaces before you can delete the policy map.
[no]	Deletes the policy map.

Creating a class map

Create a class map within the policy map to classify the traffic as desired. The class map describes the classification attributes and actions. A policy map can have multiple class maps.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To create the class map, enter:

   ```
   [no] class-map <class-name> <parent-class-name>
   ```

Table 18: Variable definitions

Variable	Value
<class-name>	The name of the class map to be created.
<parent-class-name>	Parent class name for the class map. "root" is the pre-created class map that is the root class map for the policy map. All traffic on an interface is matched to the root class.
[no]	Deletes the class-map from the policy map.

Configuring classification attributes for a class map

Configure classification parameters to specify the attributes that the packets must match in order to be classified into the specified class map.

To create a default class map, enter default as the match value.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure the classification keys, enter:

   ```
   [no] match
   [packet-class <packet-classvalue/string>]
   [src-ip <src-ip>[/<netmask>]]
   [dest-ip <destip>[/<netmask>]]
   [port <1-65535>]
   [step-size <1-65535>]
   [vlan-id <1-4095>]
   [dscp <string/0-63>]
   [precedence <0-7>]
   [tos <0-255>]
   [user-priority <0-7>]
   [protocol <value/string>]
   [exp <0-7>]
   [src-ipv6 <src-ipv6>[/prefixlen]]
   [dest-ipv6 <dest-ipv6>[/<prefixlen <1-128>]]
   [traffic-class <0-255>]
   [flow-label <0-1048575>]
   [mpls-label <0-1048575>]
   ```

Table 19: Variable definitions

Variable	Value
[packet-class <packet-classvalue/string>]	Match protocol type in layer 2 header or 'default' (E.g. packet-class IPV4). Valid: {IPV4, IPV6, MPLS} -or- <hex-value (0001-FFFF)>.
[src-ip <src-ip>[/ <netmask>]]	Match source IP address. Valid: IP address/range/subnet or 'default' Range of subnets are not allowed. Always defaults to 32 (255.255.255.255).
[dest-ip <destip>[/ <netmask>]]	Match destination IP address. Valid: IP address/range/ subnet or 'default'. Range of subnets are not allowed. Always defaults to 32 (255.255.255.255).
[port <1-65535>]	Match TCP/UDP ports. Valid: port number/range or 'default'. (1-65535)
[step-size <1-65535>]	Step increment (for example, **`match port 21-1000 step-size 2`**). Defaults to 1. Valid: 1-65535.
[vlan-id <1-4095>]	VLAN ID/range or 'default' . Valid: 1-4095.
[dscp <string/0-63>]	Match DiffServ codepoints. DSCP/range or 'default'. Valid: 0-63, ef, af and cs code point
[precedence <0-7>]	Match precedence in IP header. Precedence value/range or 'default'. Valid: 0-7
[tos <0-255>]	IP ToS field value/range or 'default' (e.g. tos 5-8). Valid: 0-255.
[user-priority <0-7>]	Match 802.1p priority. Priority value/range or 'default'. Valid: 0-7
[protocol <value/string>]	Match Protocol type in IP header. Protocol id or 'default'. Valid: {UDP TCP OSPF MPLS} -or- value (1-255)
[exp <0-7>]	Match MPLS EXP bits. EXP value/range or 'default'. Valid : 0-7
[src-ipv6 <src-ipv6>[/ prefixlen]]	Match IPv6 source address. IPv6 address/prefix or 'default'. The prefix, if not specified defaults to 128.
[dest-ipv6 <dest-ipv6>[/ <prefixlen <1-128>]]	Match IPv6 destination address. IPv6 address/prefix or 'default'. The prefix, if not specified defaults to 128.
[traffic-class <0-255>]	Match Traffic Class field of IPv6. Value/range or 'default'. Valid: 0-255
[flow-label <0-1048575>]	IPv6 Flow Label value/range or 'default'. Valid: 0- 1048575
[mpls-label <0-1048575>]	MPLS Label value/range or 'default' . Valid : 0-1048575
[no]	Removes the classification keys from the class-map.

Cloning policy maps

Clone a policy-map to copy an existing configuration to another policy-map name. The new policy-map can then be edited to create a new configuration.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To clone the policy map, enter:

   ```
   clone-policy-map <new-policy-name> <existing-policy-name>
   ```

Table 20: Variable definitions

Variable	Value
<new-policy-name>	Name of the policy map to be created.
<existing-policy-name>	The policy-map to be copied.

Applying the policy map to an interface

After you have configured a policy map, you can assign it to one or more interfaces.

The policy map can be associated with an interface in the inbound or outbound direction. When you associate a policy map with an interface, an instance of the policy map is instantiated over the interface in the specified direction. If you change the policy map after applying it to the interface, the changes are applied across all associated interfaces.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface <interface-type> <interface-name>
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To clone the policy map to the interface, enter:

   ```
   [no] service-policy {input | output} <policy-name>
   ```

Table 21: Variable definitions

Variable	Value
<interface-type> <interface-name>	Interface type can be Ethernet, bundle, tunnel, vlan or crypto.
{input | output}	Specifies whether the policy map is applied in the inbound direction or outbound direction.
<policy-name>	The policy map name that has to be applied.
[no]	Removes the policy map from that interface.

Mapping a priority queue to a class map

Use this procedure to map a Strict Priority Queuing (SPQ) to a class map.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select qos chassis, enter:

   ```
   qos chassis
   ```

3. To specify the policy map, enter:

   ```
   policy-map <policy-map-name>
   ```

4. To specify the class map, enter:

   ```
   class-map <class-map-name>
   ```

5. To specify which spq queue, enter:

   ```
   assign-queue <1 - 8>
   ```

6. To exit the class map configuration mode, enter:

   ```
   exit
   ```

7. To exit the policy map configuration mode, enter:

   ```
   exit
   ```

Table 22: Variable definitions

Variable	Value
assign-queue <1 - 8>	Specifies the assigned output queue. Values range from 1 to 8.

<class-map-name>	Specifies the class map to configure under the policy map.
<policy-map-name>	Specifies the name of the policy map.

Configuring QoS over Frame Relay

To configure QoS over Frame Relay, apply a policy map to the PVC:

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface [ bundle <bundle-name>]
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To clone the policy map to the PVC, enter:

   ```
   [no] service-policy {input | output} <policy-name> pvc <PVC>
   ```

Table 23: Variable definitions

Variable	Value
{input | output}	Specifies whether the policy map is applied in the inbound direction or outbound direction.
<policy-name>	The policy map name that has to be applied.
[no]	Removes the policy map from the specified PVC.
<PVC>	Specifies the PVC identifier: 16-1022.

Configuring CoS marking

Configure CoS marking for a class map to mark the header fields of the matching class packets.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To configure marking, enter:

```
[no] mark {[dscp <string/0-63>] | [precedence <0-7>] | [user-
priority <0-7>] | [exp <0-7>]}
```

Table 24: Variable definitions

Variable	Value
[dscp <string/0-63>]	Mark packets with DiffServ Codepoint. Valid: 0-63, ef, af and cs code points.
[precedence <0-7>]	Mark packets with IP precedence. Valid: 0-7.
[user-priority <0-7>]	Mark packets with 802.1p priority. Valid: 0-7.
[exp <0-7>]	Mark packets with EXP value. Valid: 0-7.
[no]	Deletes the specified marking configuration for the class.

Configuring traffic policing

Configuring Policing using Single Rate Three Color Marker

Configure srTCM related parameters to apply single-rate policing to a class map in the inbound or outbound direction.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To select the policy map, enter:

Comments? infodev@avaya.com

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To select policing, enter:

```
police
```

6. To configure srTCM, enter:

```
[no] srtcm
[cir <1-1000000>]
[cir-percentage <cir%>]
[cbs <1-5000000>]
[cbs-time <100-5000>]
[ebs <1-5000000>]
[ebs-time <100-5000>]
[green-action {permit | mark-dscp-<value> | drop}]
[yellow-action {permit | mark-dscp-< value> | drop}]
[red-action {permit | markdscp-<value> | drop}]
```

Table 25: Variable definitions

Variable	Value
[cir <1-1000000>]	Committed rate in Kbps. Valid: 1 – 1000000.
[cir-percentage <cir%>]	Committed rate in percentage.
[cbs <1 – 5000000>]	Committed burst size in Kbits. The valid range corresponding to CIR should be >= (100 msec * CIR)/1000 and <= (5000 msec * CIR)/1000.
[cbs-time <100-5000>]	Committed burst size in milliseconds. The valid range is 100 to 5000. Default value is 1000 ms.
[ebs <1 – 5000000>]	Excess burst size in Kbits. The valid range corresponding to CIR should be >= (100 msec * CIR)/1000 and <= (5000 msec * CIR)/1000.
[ebs-time <100-5000>]	Excess burst size in milliseconds. The valid range is 100 to 5000. Default value is 1000 ms.
[green-action {permit | mark-dscp-<value> | drop}]	Action to be taken on green packet. Valid permit, mark-dscp-<value> or drop.
[yellow-action {permit | mark-dscp-< value> | drop}]	Action to be taken on yellow packet. Valid permit, mark-dscp-<value> or drop

Variable	Value
[red-action {permit | markdscp-<value> | drop}]	Action to be taken on red packet. Valid permit, mark-dscp-<value> or drop
[no]	Deletes the specified policing configuration for the class.

Configuring Policing using Two Rate Three Color Marker

Configure trTCM related parameters to apply two-rate policing to a class map in the inbound or outbound direction.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To select policing, enter:

   ```
   police
   ```

6. To configure trTCM, enter:

   ```
   [no] trtcm
   [cir <cir>]
   [cir-percentage <cir%>]
   [pir <pir>]
   [pir-percentage <pir%>]
   [cbs <1 - 5000000>]
   [cbs-time <100-5000>]
   [pbs <1 - 5000000>]
   [pbs-time <100-5000>]
   [green-action {permit | mark-dscp-<value> | drop}]
   [yellow-action {permit | mark-dscp-<value> | drop}]
   ```

```
[red-action {permit | markdscp-<value> | drop}]
```

Table 26: Variable definitions

Variable	Value
[cir <cir>]	Committed rate in Kbps. Valid: 1 – 1000000.
[cir-percentage <cir%>]	Committed rate in percentage.
[pir <pir>]	Peak rate in Kbps. Valid: 1 – 1000000.
[pir-percentage <pir%>]	Peak rate in percentage.
[cbs <1 - 5000000>]	Committed burst size in Kbits. The valid range corresponding to CIR should be >= (100 msec * CIR)/1000 and <= (5000 msec * CIR)/1000.
[cbs-time <100-5000>]	Committed burst size in milliseconds. The valid range is 100 to 5000. Default value is 1000 ms.
[pbs <1 - 5000000>]	Peak burst size in Kbits. The valid range corresponding to PIR should be >= (100 msec * PIR)/1000 and <= (5000 msec * PIR)/1000.
[pbs-time <100-5000>]	Peak burst size in milliseconds. The valid range is 100 to 5000. Default value is 1000 ms.
[green-action {permit | mark-dscp-<value> | drop}]	Action to be taken on green packet. Valid permit, mark-dscp-<value> or drop.
[yellow-action {permit | mark-dscp-<value> | drop}]	Action to be taken on yellow packet. Valid permit, mark-dscp-<value> or drop
[red-action {permit | markdscp-<value> | drop}]	Action to be taken on red packet. Valid permit, mark-dscp-<value> or drop
[no]	Deletes the specified policing configuration for the class.

Configuring color aware mode for srTCM and trTCM

Enable srTCM or trTCM to run in color aware mode for the class map to allow consideration of previous DSCP markings.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To select policing, enter:

```
police
```

6. To configure color-aware mode, enter:

```
[no] color-aware
```

Table 27: Variable definitions

Variable	Value
[no]	Disables color aware mode.

Configuring policy-based redirect

Configure policy-based redirect to redirect traffic from a specified class to a particular interface.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To configure PBR, enter:

```
[no] pbr-redirect {[nexthop <ip-address>] | [lsp <lsp-name>]
| [interface <interface name>]}
```

Table 28: Variable definitions

Variable	Value
[nexthop <ip-address>]	Nexthop IPv4/IPv6 address.
[lsp <lsp-name>]	MPLS LSP interface name.

Variable	Value
[interface <interface name>]	Egress interface name.
[no]	Resets the values to null. It will disable the feature as well.

Configuring congestion control and avoidance

Configuring WRED or DS-RED parameters for a class

Specify WRED or DS-RED parameters on a class to configure congestion avoidance on outbound flows.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure WRED or DS-RED

   ```
   [no] red

   {[minth-green <percent>] [maxth-green <percent>] [mpdgreen
   <1-15>] [minth-yellow <percent>] [maxth-yellow <percent>]
   [mpd-yellow <1-15>] [minth-red <percent>] [maxth-red
   <percent>] [mpd-red <1-15>] } |

   {[dscp <0-63>] [minth <percent>] [maxth <percent>]}

   [mpd <1-15>]
   ```

Table 29: Variable definitions

Variable	Value
[minth-green <percent>]	Minimum Threshold for green (WRED). Valid Range: 1-100.

Variable	Value
[minth-yellow <percent>]	Minimum Threshold for yellow (WRED). Valid Range: 1-100.
[minth-red <percent>]	Minimum Threshold for red (WRED). Valid Range: 1- 100.
[maxth-green <percent>]	Maximum Threshold for green (WRED). Valid Range: 1-100.
[maxth-yellow <percent>]	Maximum Threshold for yellow (WRED). Valid Range: 1-100.
[maxth-red <percent>]	Maximum Threshold for red (WRED). Valid Range: 1- 100.
[mpdgreen <1-15>]	Mark Probability Denominator (as power of 2) for green(WRED). Valid Range: 1-15.
[mpd-yellow <1-15>]	Mark Probability Denominator (as power of 2) for yellow(WRED). Valid Range: 1-15.
[mpd-red <1-15>]	Mark Probability Denominator (as power of 2) for red(WRED). Valid Range: 1-15.
[minth <percent>]	Minimum Threshold for DS-RED. Valid Range: 1-100.
[maxth <percent>]	Maximum Threshold for DS-RED. Valid Range: 1-100.
[dscp <0-63>]	DiffServ Codepoint for DS-RED. Valid Range: 0-63/string.
[mpd <1-15>]	Mark Probability Denominator (as power of 2) for DS-RED. Valid Range: 1-15.
[no]	Resets the WRED parameters to its default values, which depends on the interface, and removes the DS-RED parameters if specified.

Configuring EWF for a class

Configure the Exponential Weight Factor on this class for average queue size calculation

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To configure EWF, enter:

```
[no] ewf <1-15>
```

Table 30: Variable definitions

Variable	Value
<1-15>	Specifies the EWF value.
[no]	Removes the configuration.

Enabling WRED or DS-RED for a class

Enable WRED or DS-RED on a class to enable congestion avoidance on outbound flows.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To enable WRED or DS-RED

```
[no] enable-red {wred | ds-red}
```

Table 31: Variable definitions

Variable	Value
{wred | ds-red}	RED type to be enabled, either WRED or DS-RED.
[no]	Disables the corresponding RED type.

Configuring WRED or DS-RED parameters for an interface

Specify WRED or DS-RED parameters on an interface to configure congestion avoidance on outbound flows.

At the interface level, you can only enable WRED or DS-RED for WAN interfaces.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface [ bundle <bundle-name> ]
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To configure WRED or DS-RED

   ```
   [no] red

   [pvc <1-4096>]

   {[minth-green <percent>] [maxth-green <percent>] [mpdgreen
   <1-15>] [minth-yellow <percent>] [maxth-yellow <percent>]
   [mpd-yellow <1-15>] [minth-red <percent>] [maxth-red
   <percent>] [mpd-red <1-15>] } |

   {[dscp <0-63>] [minth <percent>] [maxth <percent>]}

   [mpd <1-15>]
   ```

Table 32: Variable definitions

Variable	Value
[pvc <1-4096>]	Specifies the PVC ID. Not required for non-Frame Relay interfaces.
[minth-green <percent>]	Minimum Threshold for green (WRED). Valid Range: 1-100.
[minth-yellow <percent>]	Minimum Threshold for yellow (WRED). Valid Range: 1-100.
[minth-red <percent>]	Minimum Threshold for red (WRED). Valid Range: 1- 100.
[maxth-green <percent>]	Maximum Threshold for green (WRED). Valid Range: 1-100.
[maxth-yellow <percent>]	Maximum Threshold for yellow (WRED). Valid Range: 1-100.
[maxth-red <percent>]	Maximum Threshold for red (WRED). Valid Range: 1- 100.
[mpdgreen <1-15>]	Mark Probability Denominator (as power of 2) for green(WRED). Valid Range: 1-15.
[mpd-yellow <1-15>]	Mark Probability Denominator (as power of 2) for yellow(WRED). Valid Range: 1-15.

Variable	Value
[mpd-red <1-15>]	Mark Probability Denominator (as power of 2) for red(WRED). Valid Range: 1-15.
[minth <percent>]	Minimum Threshold for DS-RED. Valid Range: 1-100.
[maxth <percent>]	Maximum Threshold for DS-RED. Valid Range: 1-100.
[dscp <0-63>]	DiffServ Codepoint for DS-RED. Valid Range: 0-63/string.
[mpd <1-15>]	Mark Probability Denominator (as power of 2) for DS-RED. Valid Range: 1-15.
[no]	Resets the WRED parameters to its default values, which depends on the interface, and removes the DS-RED parameters if specified.

Configuring EWF for an interface

Configure the Exponential Weight Factor on this class for average queue size calculation.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface [ bundle <bundle-name>]
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To configure EWF, enter:

   ```
   [no] ewf weighting-factor <1-15> [pvc <dlci>]
   ```

Table 33: Variable definitions

Variable	Value
<1-15>	Specifies the EWF value.
[no]	Removes the configuration.
[pvc <dlci>]	PVC identifier. Optional for non-frame relay interfaces

Enabling WRED or DS-RED on the interface

Enable WRED or DS-RED on an interface to enable congestion avoidance on outbound flows.

At the interface level, you can only enable WRED or DS-RED for WAN interfaces.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface [ bundle <bundle-name>]
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To enable WRED or DS-RED

   ```
   [no] enable-red {wred | ds-red}
   ```

Table 34: Variable definitions

Variable	Value
{wred | ds-red}	RED type to be enabled, either WRED or DS-RED.
[no]	Disables the corresponding RED type.

Configuring queueing and scheduling

Configuring CBQ shaping parameters

Configure CBQ parameters to enable shaping on outbound queues.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure CBQ, enter:

   ```
   [no] cbq [cr-percent <value>] [pr-percent <value>] [priority
   <1-8>]
   ```

Table 35: Variable definitions

Variable	Value
[cr-percent <value>]	Committed Rate as percentage of interface bandwidth. Meaningful only for leaf classes.
[pr-percent <value>]	Peak Rate as percentage of interface bandwidth.
[priority <1-8>]	Scheduling priority. Meaningful only for leaf classes
[no]	Clears the CBQ parameters.

Configuring interface shaping parameters

Configure shaping parameters for outbound traffic on a chassis Ethernet interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To configure shaping, enter:

   ```
   shaping <value>
   ```

Table 36: Variable definitions

Variable	Value
<value>	Value of the interface bandwidth in kbps; Default value is 50000.

Configuring committed rate for priority queue on an Ethernet interface

Use this procedure to configure SPQ on an Ethernet interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the ethernet, enter:

   ```
   interface ethernet <slot/port number>
   ```

3. To specify QoS chassis section, enter:

   ```
   qos chassis
   ```

4. To specify the SPQ Queue number, enter:

   ```
   queue <1--8>
   ```

5. To specify committed rate percentage, enter:

   ```
   shape <shaping_rate>
   ```

6. To exit the queue configuration mode, enter:

   ```
   exit
   ```

Table 37: Variable definitions

Variable	Value
queue <1--8>	Specifies the output queue. Values range from 1 to 8.
shape <shaping_rate>	Specifies the shaping rate as a percentage of the interface bandwidth. Values range from 0 to 100.

Configuring RED for priority queue on Bundle interface

Use this procedure to configure WRED or DS-RED on a bundle interface.

Procedure steps

1. To enter configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface bundle, enter:

```
interface bundle <bundle-name>
```

3. To specify link info, enter:

```
link <interface-type> <interface number>
```

4. To specify bundle encapsulation, enter:

```
encap <encap-type>
```

5. To specify QoS chassis section, enter:

```
qos chassis
```

6. To specify the SPQ Queue number, enter:

```
queue <1--8>
```

7. To specify RED thresholds for each drop precedence, enter:

```
[no] red [pvc <1-4096>] [minth-red <percent>] [minth-yellow
<percent>] [minth-green <percent>] [maxth-red <percent>]
[maxth-yellow <percent>] [maxth-green <percent>] [mpd-red
<1-15>] [mpd-yellow <1-15>] [mpd-green <1-15>] [minth
<percent>] [maxth <percent>] [mpd <1-15>] [dscp <0-63>]
```

8. To enable RED on the queue, enter:

```
enable-red
```

9. Toexit the queue configuration mode, enter:

```
exit
```

Table 38: Variable definitions

Variable	Value
[dscp <0-63>]	DiffServ Codepoint for DS-RED. Valid Range: 0-63/string.
enable-red	Enables WRED or DS-RED on an interface to enable congestion avoidance on outbound flows.
encap <encap-type>	Specifies encapsulation type for the interface bundle. Values include: • ppp • hdlc • frelay • mlppp • mfr • atm
interface bundle <bundle-name>	Identifies the interface bundle.

Variable	Value
link <type> <interface number>	Specifies the link type interface type and interface number to associate with the link type. Link types include: • bri • dialer • e1 • pri_e1 • serial • t3
[minth-red <percent>]	Minimum Threshold for red (WRED). Valid Range: 1- 100.
[minth-yellow <percent>]	Minimum Threshold for yellow (WRED). Valid Range: 1-100.
[minth-green <percent>]	Minimum Threshold for green (WRED). Valid Range: 1-100.
[maxth-red <percent>]	Maximum Threshold for red (WRED). Valid Range: 1- 100.
[maxth-yellow <percent>]	Maximum Threshold for yellow (WRED). Valid Range: 1-100.
[maxth-green <percent>]	Maximum Threshold for green (WRED). Valid Range: 1-100.
[mpd-red <1-15>]	Mark Probability Denominator (as power of 2) for red(WRED). Valid Range: 1-15.
[mpd-yellow <1-15>]	Mark Probability Denominator (as power of 2) for yellow(WRED). Valid Range: 1-15.
[mpd-green <1-15>]	Mark Probability Denominator (as power of 2) for green(WRED). Valid Range: 1-15.
[minth <percent>]	Minimum Threshold for DS-RED. Valid Range: 1-100.
[maxth <percent>]	Maximum Threshold for DS-RED. Valid Range: 1-100.
[mpd <1-15>]	Mark Probability Denominator (as power of 2) for DS-RED. Valid Range: 1-15.
[no]	Resets the WRED parameters to their default values, which depends on the interface, and removes the DS-RED parameters if specified.

Variable	Value
[pvc <1-4096>]	Specifies the PVC ID. Not required for non-Frame Relay interfaces.
queue <1--8>	Specifies the queue value for the interface bundle. Values range from 1 to 8.

Configuring SPQ on an Ethernet interface

Use this procedure to enable or disable SPQ on an Ethernet interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the ethernet, enter:

   ```
   interface ethernet <slot/port number>
   ```

3. To specify QoS chassis section, enter:

   ```
   qos chassis
   ```

4. To enable SPQ, enter:

   ```
   [no] enable spq output
   ```

Table 39: Variable definitions

Variable	Value
[no]	Disables SPQ on the Ethernet interface.
<slot/port number>	Specifies the chassis slot and port number for the Ethernet interface.

Configuring global MPLS DSCP to EXP markings

Configure MPLS EXP to DSCP markings.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To configure EXP to DSCP markings, enter:

```
dscp-exp-cos-map <0-63> exp <0-7>
```

Table 40: Variable definitions

Variable	Value
<0-63>	DSCP value. You can also specify as ef, af, and cs code points.
<0-7>	MPLS EXP value.

Configuring interface MPLS DSCP to EXP markings

Configure MPLS EXP to DSCP markings for an interface.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select a chassis interface, enter:

```
interface [ bundle <bundle-name> | ethernet <0/1-0/4>]
```

3. To specify Chassis QoS, enter:

```
qos chassis
```

4. To configure EXP to DSCP markings, enter:

```
dscp-exp-cos-map <0-63> exp <0-7>
```

Table 41: Variable definitions

Variable	Value
<0-63>	DSCP value. You can also specify as ef, af, and cs code points.
<0-7>	MPLS EXP value.

Configuring global MPLS EXP to DSCP markings

Configure MPLS EXP to DSCP markings.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To configure EXP to DSCP markings, enter:

   ```
   exp-dscp-cos-map <0-7> dscp <0-63>
   ```

Table 42: Variable definitions

Variable	Value
<0-63>	DSCP value. You can also specify as ef, af, and cs code points.
<0-7>	MPLS EXP value.

Configuring interface MPLS EXP to DSCP markings

Configure MPLS EXP to DSCP markings.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface [ bundle <bundle-name> | ethernet <0/1-0/4>]
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To configure EXP to DSCP markings, enter:

   ```
   exp-dscp-cos-map <0-7> dscp <0-63>
   ```

Table 43: Variable definitions

Variable	Value
<0-63>	DSCP value. You can also specify as ef, af, and cs code points.
<0-7>	MPLS EXP value.

Configuring Buffer Management

Configure buffer management to set the maximum buffer limit for the class queue.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure buffer management

   ```
   [no] excess-queue-buffers <max-allowed-buffer-percentage>
   ```

Table 44: Variable definitions

Variable	Value
<max-allowed-buffer-percentage>	Excess maximum queue buffer limit, in percentage.
[no]	Resets the values based on the class parameters.

Enabling QoS features on an interface

To enable the features that are specified in the policy map that is applied on an interface, you must explicitly enable the specified QoS features on the interface .

Important:

On Frame Relay bundles, whenever a PVC is created, it uses all the available or unassigned bandwidth. When CBQ is enabled either manually or using auto QoS, subsequent PVC creation fails because there is no available bandwidth. To avoid this scenario, you must change the bandwidth for created PVCs such that there is always non-zero bandwidth available.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface <interface-type> <interface-name>
   ```

3. To specify Chassis QoS, enter:

   ```
   qos chassis
   ```

4. To enable the desired QoS features, enter:

   ```
   [no] enable [cbq | policing | monitoring | pbr] {inbound | outbound}
   ```

Table 45: Variable definitions

Variable	Value
cbq	Enables CBQ on the interface. CBQ, policing, and monitoring are mutually-exclusive features in either direction on the interface.
{inbound | outbound}	Specifies the traffic direction in which the feature applies.
policing	Enables policing on the interface.
monitoring	Enables monitoring on the interface.
pbr	Enables policy-based redirect on the interface.
[no]	Disables the specified feature.

Configuring Statistics

Configuring the sample interval for statistics

Configure the sample interval for statistics.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To specify historical statistics, enter:

```
historical-stats
```

4. To configure the sample interval, enter:

```
sample-interval <interval>
```

Table 46: Variable definitions

Variable	Value
<interval>	Sample Interval length (5, 10 or 15 minutes).

Configuring FTP parameters for upload of statistics

Configure FTP parameters for the upload of statistics:

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To specify historical statistics, enter:

```
historical-stats
```

4. To configure FTP parameters, enter:

```
ftp-parameters
```

The system prompts you to enter values for the primary FTP server IP address, secondary FTP server IP address, user name, and password.

Configuring file parameters and interval for the upload of statistics

Configure the file name and interval for upload of statistics to an FTP server.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To specify historical statistics, enter:

   ```
   historical-stats
   ```

4. To configure upload parameters, enter:

   ```
   [no] upload [interval <value>] [file-id <name>]
   ```

Table 47: Variable definitions

Variable	Value
[interval <value>]	Upload Interval (1, 2, 3 or 4 hours).
[file-id <name>]	Upload-file ID.
[no]	Disables uploading of historical statistics.

Disabling QoS globally

Enable or disable QoS for all the applicable interfaces.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify QoS, enter:

   ```
   qos
   ```

3. To disable QoS, enter:

   ```
   [no] disable-qos
   ```

Table 48: Variable definitions

Variable	Value
[no]	Re-enables QoS at global level.

Disabling QoS on an interface

Enable or disable QoS for the specified interface.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select a chassis interface, enter:

```
interface <interface-type> <interface-name>
```

3. To specify QoS configuration, enter:

```
qos
```

4. To disable QoS, enter:

```
[no] disable-qos
```

Table 49: Variable definitions

Variable	Value
[no]	Re-enables QoS at interface level.

Configuring auto QoS globally

Enable auto QoS at the global level to apply auto QoS to all the applicable interfaces (chassis and module interfaces). If an interface already has a policy map applied, auto QoS is not applied to that interface.

Important:

For PPP bundles, when auto QoS is enabled, if the bundle bandwidth is less than or equal to 768Kbps, the LFI feature is enabled automatically. The feature is disabled when auto QoS is removed.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify QoS configuration, enter:

```
qos
```

3. To configure auto QoS, enter:

```
[no] enable-auto-qos
```

Table 50: Variable definitions

Variable	Value
[no]	Removes the auto QoS configuration.

Configuring auto QoS on an interface

Enable auto QoS at the interface level to apply auto QoS to the interface. If the interface already has a policy map applied, auto QoS is not applied.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select a chassis interface, enter:

   ```
   interface <interface-type> <interface-name>
   ```

3. To specify Chassis QoS, enter:

   ```
   qos
   ```

4. To configure auto QoS, enter:

   ```
   [no] enable-auto-qos
   ```

Table 51: Variable definitions

Variable	Value
<interface-type> <interface-name>	Specifies the interface to configure. Valid values can be one of the following: bundle <bundle-name> ethernet <0/1-0/4> vlan <2-4094> tunnel <tunnel-name> crypto
[no]	Removes the auto QoS configuration from the interface.

Displaying QoS configuration and statistics

Displaying policy maps and class maps

Display policy map and class map information.

Procedure steps

To display policy maps and class map configuration, enter:

```
show qos chassis policy-map [<policy-map-name>] [class-map
[<class-map-name>]] [detail]
```

Table 52: Variable definitions

Variable	Value
[<policy-map-name>]	Name of policy map to display. String(1-19). If you do not specify a policy map name, the system displays all the configured policy maps.
[<class-map-name>]	Name of class map to display. String(1-19). If you do not specify a class map name, the system displays all the configured class maps for the specified policy map.
[detail]	Displays detailed information.

Displaying policy maps and class maps for an interface

Display service policy map information for the specified interface. If a class map is specified then the output displays information about that class map only.

Procedure steps

To display interface policy maps and class maps, enter:

```
show qos chassis <interface-type> <interface-name> [pvc <pvc-
id>] [class <class-map-name>] [direction {inbound | outbound}]
```

Table 53: Variable definitions

Variable	Value
[direction {inbound | outbound}]	Specifies to display statistics for inbound or outbound direction.
<interface-type>	Specifies the interface type: ethernet, bundle, tunnel, vlan or crypto.
<interface-name>	Name of the interface.
[class-map <class-map-name>]	Class map name to display. String(1-19)
[pvc <pvc-id>]	PVC identifier for the FR interface. Mandatory parameter for FR interfaces.
[mpls]	Display MPLS LSP tunnel name information.
[lsp <lsp>]	MPLS LSP interface name.

Displaying mapping of interfaces to policy maps

Display service policy map information to view the policy maps that are applied to all applicable interfaces. The output displays in the format of interface and inbound and outbound service policy maps.

Procedure steps

To display service policy map information, enter:

```
show qos chassis service-policy <interface-type> <interface-name> [pvc <pvc-id>]
```

Table 54: Variable definitions

Variable	Value
<interface-type>	Specifies the interface type: ethernet, bundle, tunnel, vlan or crypto.
<interface-name>	Name of the interface.
[pvc <pvc-id>]	Specifies the PVC ID for Frame Relay interfaces. Range: 16-1022.

Displaying system level QoS configuration

Display the system level QoS configuration.

Procedure steps

To display the system level QoS configuration, enter:

```
show qos chassis system
```

Displaying the DSCP to EXP mappings

Display the DSCP to EXP mappings for MPLS QoS.

Procedure steps

To display the DSCP to EXP mappings, enter:

```
show qos chassis dscp-exp-cos-map
```

Displaying the EXP to DSCP mappings

Display the EXP to DSCP mappings for MPLS QoS.

Procedure steps

To display the EXP to DSCP mappings, enter:

```
show qos chassis exp-dscp-cos-map
```

Displaying the configuration for historical statistics

Execute this command to display the configuration for the collection of historical statistics.

Procedure steps

To display the configuration for historical statistics, enter:

```
show qos chassis historical-stats configuration
```

Displaying historical statistics

Display the historical statistics for class maps for the specified interface and/or specified class map..

Procedure steps

To display historical statistics, enter:

```
show qos chassis historical-stats <interface-type> <interface-
name> [pvc < pvc-id>] [class-map <class-mapname>] [direction
{inbound | outbound}]
```

Table 55: Variable definitions

Variable	Value
[direction {inbound | outbound}]	Specifies to display statistics for inbound or outbound direction.
<interface-type>	Specifies the interface type: bundle, ethernet, tunnel, or vlan, or crypto.
<interface-name>	Specifies the interface name for the bundle, ethernet, tunnel, or vlan.

Variable	Value
[pvc < pvc-id>]	PVC identifier for the FR interface. Mandatory parameter for FR interfaces.
[class-map <class-mapname>]	Class map name for which historical statistics to display. If not specified, then it will display for all class maps.

Clearing QoS statistics

Clear the QoS statistics for the specified interface and/or class map.

Execute this command to clear the statistics information for the specified interface. If class map is specified then it clear statistics for the class map only.

Procedure steps

To clear QoS statistics, enter:

```
clear qos chassis <interfacetype> <interfacename> [pvc < pvc-id>] [class <class-map-name>]
```

Table 56: Variable definitions

Variable	Value
<interfacetype>	Can be Ethernet, bundle, tunnel, vlan or crypto.
<interfacename>	Name of the interface.
[class-map <class-map-name>]	Class map name to clear the statistics. String(1-19)
[pvc < pvc-id>]	PVC identifier for the FR interface. Mandatory parameter for FR interfaces.

Displaying RED information for a bundle

Display the RED information for the specified bundle.

Procedure steps

To display RED information, enter:

```
show qos chassis red <bundle-name> [pvc < pvc-id>]
```

Table 57: Variable definitions

Variable	Value
[pvc < pvc-id>]	PVC identifier for the FR interface. Mandatory parameter for FR interfaces.
<bundle-name>	Specifies the name for the bundle.

Clearing RED information for a bundle

Clear the RED information for the specified bundle.

Procedure steps

To clear RED information, enter:

```
clear qos chassis red <bundle-name> [pvc < pvc-id>]
```

Table 58: Variable definitions

Variable	Value
[pvc < pvc-id>]	PVC identifier for the FR interface. Mandatory parameter for FR interfaces.
<bundle-name>	Specifies the name for the bundle.

Chapter 8: Ethernet Module QoS configuration

The high-level configuration steps for Ethernet Module QoS are:

1. Configure CoS mapping at L2 and interface level.
2. Create the policy map.
3. Configure the class-maps to classify the traffic.
4. Associate actions with the leaf classes.
5. Apply the policy map to the interface in the inbound direction.
6. Configure WRED.
7. Configure SPQ/WRR.
8. Configure shaping.

Important:

The GigE Large Ethernet module is oversubscribed and provides a 4:1 blocking ratio. That is, a bandwidth of one gigabit is shared by every four ports. Within a group of four ports, any QoS treatment applied to one port is also applied to the three remaining ports in the group, with the exception of the following features, which are applied independently for each port:

- QoS status
- Auto QoS status
- Interface DSCP marking
- Interface UP marking
- Service policy attached to the interface (policy map and its associated configuration)

Configuring L2-based CoS marking

Configuring the default output queue

Assign the default output queue to specify the queue assigned to all untagged packets ingressing on the interface. To force output-queue assignment for both untagged and tagged packets, you can use the **`force-all`** parameter.

Important:

The Avaya Secure Router 2330 does not support port-level queue assignment for tagged packets, so the **default-queue** command is not supported. Output queue for packets ingressing through the interface is obtained from the **user-priority-cos-map** table using the UP value obtained from either the incoming packet, if tagged, or from the interface default UP configuration (configured using the **mark-user-priority** command), if untagged.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure the default output queue for the interface, enter:

   ```
   [no] default-queue <queue-no> [force-all]
   ```

Table 59: Variable definitions

Variable	Value
<queue-no>	The output queue assigned. Valid Range: 1-8
[force-all]	Optional parameter. If specified, output-queue assignment is forced for both untagged and tagged packets. Otherwise, it is assigned only for untagged packets.
[no]	Resets the default queue to 8.

Mapping user priority to output queue and drop precedence

Map the user priority value to an output queue and drop precedence value for priority tagged packets ingressing on the interface.

Important:

The **user-priority-cos-map** command is not supported on the SR2330. The user-priority in the Layer 2 header determines the output queue.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Module QoS, enter:

```
qos module
```

3. To configure user priority mapping, enter:

```
[no] user-priority-cos-map <userpriority-value> { [queue
<queue-no> ] [drop-precedence <low | medium | high>] }
```

Table 60: Variable definitions

Variable	Value
<userpriority-value>	The user priority value of the priority tagged packet. Valid Range: 0-7
[queue <queue-no>]	The output queue assigned. Valid Range: 1-8
[drop-precedence <low | medium | high>]	The drop precedence value assigned. Valid Range: low / medium / high.
[no]	Resets the user-priority CoS mapping to default values.

Mapping output queue to a user priority value for untagged packets

Map the output queue to a user priority value for untagged packets egressing on that queue.

Important:

On the SR2330, the **queue-cos-map** command is not supported. Queue mapping is derived from the **user-priority-cos-map** table using the UP value.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Module QoS, enter:

```
qos module
```

3. To configure output queue mapping, enter:

```
[no] queue-cos-map <queue-no> <userpriority-value>
```

Table 61: Variable definitions

Variable	Value
<queue-no>	The output queue assigned. Valid Range: 1-8

Variable	Value
<userpriority-value>	The user priority value of the priority tagged packet. Valid Range: 0-7
[no]	Resets the queue and user-priority mapping to default values.

Configuring interface-based CoS marking

Configuring user priority for packets ingressing on an interface

Configure the user priority value to apply for marking of untagged packets ingressing on the interface

Important:

On the SR2330, the **mark-user-priority** command is not supported.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure user-priority marking for ingress packets, enter:

   ```
   [no] mark-user-priority <user-priority-value>
   ```

Table 62: Variable definitions

Variable	Value
<user-priority-value>	The user priority value marked for the untagged packet. Valid Range: 0-7
[no]	Removes the marking

Marking DSCP value for packets ingressing on an interface

Configure the DSCP value to apply for marking of untagged packets ingressing on the interface.

Important:

On the SR2330, the **mark-dscp** command is not supported at the interface level.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure DSCP marking for ingress packets, enter:

   ```
   [no] mark-dscp <dscp-value>
   ```

Table 63: Variable definitions

Variable	Value
<dscp-value>	The DSCP value marked for the IP packets. Valid Range: 0-63
[no]	Removes the marking

Configuring multifield traffic classification

Creating policy map

Create a policy map to allow for the creation of class maps and associated actions.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos module
```

3. To create the policy map, enter:

```
[no] policy-map <policy-name> [force]
```

Table 64: Variable definitions

Variable	Value
<policy-name>	Specifies the name of the policy map.
[force]	When specified with the **no** parameter, forces the removal of the specified policy map from all applicable interfaces before deleting the policy map. If you do not specify this parameter, you must first manually remove the policy map from all applicable interfaces before you can delete the policy map.
[no]	Deletes the policy map.

Creating a class map

Create a class map within the policy map to classify the traffic as desired. The class map describes the classification attributes and actions. A policy map can have multiple class maps.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos chassis
```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To create the class map, enter:

```
[no] class-map <class-name>
```

Table 65: Variable definitions

Variable	Value
<class-name>	The name of the class map to be created.
[no]	Deletes the class-map from the policy map.

Configuring IPv4 matching attributes for a class map

Configure the match criteria to classify IPv4 packets for a class map.

The match criteria "any" cannot be used along with other criteria in a matching rule. However other fields can be used together in a matching rule.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify module QoS, enter:

   ```
   qos module
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure the classification keys, enter:

   ```
   [no] match ipv4
   [any]
   [src-mac <src-mac-address>]
   [dest-mac <dest-mac-address>]
   [ether-type <ethernet-type-value>]
   [vlan-id <vlan-id-value>]
   [user-priority <user-priorityvalue>]
   [src-address <src-ipv4-address>]
   [dest-address <dest-ipv4-address>]
   [protocol <protocol-value>]
   [src-port <src-port-value>]
   [dest-port <dest-port-value>]
   [dscp <dscp-value>]
   [tos <tos-value>]
   [precedence <precedencevalue>]
   ```

Table 66: Variable definitions

Variable	Value
[any]	Wildcard rule to match any IPv4 packet
[src-mac <src-mac-address>]	Source MAC address in xx:xx:xx:xx:xx:xx format
[dest-mac <dest-mac-address>]	Destination MAC address in xx:xx:xx:xx:xx:xx format
[ether-type <ethernet-type-value>]	Ethernet type value in hexa-decimal notation
[vlan-id <vlan-id-value>]	VLAN Identifier value. Valid Ranges: 0-4095
[user-priority <user-priorityvalue>]	User Priority value. Valid Ranges: 0-7
[src-address <src-ipv4-address>]	Source IPv4 address
[dest-address <dest-ipv4-address>]	Destination IPv4 address
[protocol <protocol-value>]	IP protocol value
[src-port <src-port-value>]	UDP/TCP source port value
[dest-port <dest-port-value>]	UDP/TCP source port value
[dscp <dscp-value>]	IP DSCP value. Valid Range: 0-63
[tos <tos-value>]	IP TOS value. Valid Range: 0-255
[precedence <precedencevalue>]	IP Precedence value. Valid Range: 0-7

Configuring IPv6 matching attributes for a class map

Configure the match criteria to classify IPv6 packets for a class map.

The match criteria "any" cannot be used along with other criteria in a matching rule. However other fields can be used together in a matching rule.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify module QoS, enter:

   ```
   qos module
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure the classification keys, enter:

   ```
   [no] match ipv6
   [any]
   [vlan-id <vlan-id-value>]
   [userpriority <user-priority-value>]
   [src-address <src-ipv6- address>]
   [dest-address <dest-ipv6-address>]
   [protocol <protocol-value>]
   [dscp < dscp-value>]
   [traffic-class <traffic-value>]
   [flow-label <flow-label-value>]
   ```

Table 67: Variable definitions

Variable	Value
[any]	Wildcard rule to match any IPv6 packet
[vlan-id <vlan-id-value>]	VLAN Identifier value. Valid Range: 0-4095
[userpriority <user-priority-value>]	User Priority value. Valid Range: 0-7
[src-address <src-ipv6-address>]	Source IPv6 address
[dest-address <dest-ipv6-address>]	Destination IPv6 address
[protocol <protocol-value>]	IP protocol value
[dscp <dscp-value>]	IP DSCP value. Valid Range: 0-63
[traffic-class <traffic-value>]	IP TC value. Valid Range: 0-255
[flow-label <flow-label-value>]	Flow label value. Valid Range: 0- 1048575

Configuring non-IP matching attributes for a class map

Create or modify a class map that can be used for matching packets to a specified class

The match criteria "any" cannot be used along with other criteria in a matching rule. However other fields can be used together in a matching rule.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify module QoS, enter:

   ```
   qos module
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To configure the classification keys, enter:

   ```
   [no] match non-ip {any | [src-mac <mac-address>] [dest-mac
   <mac-address>] [ether-type <ethernet-type-value>] [vlan-id
   <vlan-id-value>] [user-priority <user-priority-value>]}
   ```

Table 68: Variable definitions

Variable	Value
any	Wildcard rule to match any non-IP packet
[src-mac <src-mac-address>]	Source MAC address in xx:xx:xx:xx:xx:xx format
[dest-mac <dest-mac-address>]	Destination MAC address in xx:xx:xx:xx:xx:xx format
[ether-type <ethernet-type-value>]	Ethernet type value in hexa-decimal notation
[vlan-id <vlan-id-value>]	VLAN Identifier value. Valid Ranges: 0-4095
[user-priority <user-priorityvalue>]	User Priority value. Valid Ranges: 0-7

Assigning a queue for a class map

Assign the queue for all the packets matching the classification criteria of the class.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify module QoS, enter:

   ```
   qos module
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To assign a queue for the class, enter:

   ```
   [no] assign-queue <queue-no>
   ```

Table 69: Variable definitions

Variable	Value
<queue-no>	The output queue assigned. Valid Range: 1-8
[no]	Removes the queue assignment for the class.

Assigning a drop precedence for a class map

Assign the drop precedence for all the packets matching the classification criteria of the class.

Important:

Ethernet Module QoS in SR2330 only supports congestion limits for RED and YELLOW colored packets. For GREEN colored packets, the threshold is always the maximum descriptor limit. As a result, the values entered for **`low`** Drop Precedence (that is, GREEN colored packets) in the **`drop-curve`** command are ignored.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify module QoS, enter:

```
qos module
```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To assign a drop precedence for the class, enter:

```
[no] assign-drop-precedence {low | medium | high}
```

Table 70: Variable definitions

Variable	Value
[no]	Removes the drop precedence assignment from the class.

Marking a user priority for a class map

Mark the user priority value for all the packets matching the classification criteria of the class.

Important:

On the SR2330, the **mark-user-priority** command is not supported.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify module QoS, enter:

```
qos module
```

3. To select the policy map, enter:

```
policy-map <policy-name>
```

4. To select the class map, enter:

```
class-map <class-name>
```

5. To assign a queue for the class, enter:

```
[no] mark-user-priority <0-7>
```

Table 71: Variable definitions

Variable	Value
<0-7>	The user priority value marked for the priority tagged packet.
[no]	Removes the user-priority marking for the class.

Marking DSCP for a class map

Mark the DSCP value for all the packets matching the classification criteria of the class.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify module QoS, enter:

   ```
   qos module
   ```

3. To select the policy map, enter:

   ```
   policy-map <policy-name>
   ```

4. To select the class map, enter:

   ```
   class-map <class-name>
   ```

5. To assign a queue for the class, enter:

   ```
   [no] mark-dscp <0-63>
   ```

Table 72: Variable definitions

Variable	Value
<0-63>	The DSCP value marked for the IP packets.
[no]	Removes the DSCP marking for all the packets classified into the class.

Applying the policy map to an interface

After you have configured a policy map, you can assign it to one or more interfaces. With Ethernet module QoS, you can only apply the policy map in the inbound direction.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select a chassis interface, enter:

```
interface ethernet <slot/port>
```

3. To specify Chassis QoS, enter:

```
qos module
```

4. To clone the policy map to the interface, enter:

```
[no] service-policy input <policy-name>
```

Table 73: Variable definitions

Variable	Value
<policy-name>	Name of the policy map. Name can be up to 255 alphanumeric characters
[no]	Removes the policy map from that interface.

Cloning policy maps

Clone a policy-map to copy an existing configuration to another policy-map name. The new policy-map can then be edited to create a new configuration.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Chassis QoS, enter:

```
qos module
```

3. To clone the policy map, enter:

```
clone-policy-map <new-policy-name> <existing-policy-name>
```

Table 74: Variable definitions

Variable	Value
<new-policy-name>	Name of the new policy map to be created.
<existing-policy-name>	The policy map to be copied.

Configuring traffic policing

Configuring policing-based CoS mappings for non-IP packets

Configure policing based CoS mappings for non-IP packets. You can configure user-priority and drop-precedence marking for the packets based on the conformance level and incoming user-priority.

You must specify at least one of the optional parameters when configuring the CoS map.

Important:

On the Avaya Secure Router 2330, policing-based CoS marking is supported using the **`remark-cos`** command. On the Avaya Secure Router 4134, it is supported using the global **`policing-cos-map`** command.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To configure policing, enter:

   ```
   [no] policing-cos-map non-ip <user-priority-value>
   <conformance-level>

   [drop-precedence <drop-precedence-value>]

   [user-priority <user-priority-value>]
   ```

Table 75: Variable definitions

Variable	Value
<user-priority-value>	The user priority value in the priority tagged packet. Valid Range: 0-7
<conformance-level>	The conformance-level assigned by the policer to the packet. Valid Range: conform / exceed / violate
[drop-precedence <drop-precedence-value>]	Optional. The drop precedence value to be assigned. Valid Range: low / medium / high

Variable	Value
[user-priority <user-priority-value>]	Optional. The user priority value to be assigned. Valid Range: 0-7
[no]	Resets the CoS mappings for the packet

Configuring policing-based CoS mappings for IP packets

Configure policing based CoS mappings for IP packets. You can configure user-priority, DSCP, and drop-precedence marking for the packets based on the conformance level and incoming DSCP value.

One of the optional parameter must be specified while configuring the CoS map

Important:

On the SR2330, policing-based CoS marking is supported using the **remark-cos** command. On the SR4134, it is supported using the global **policing-cos-map** command.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To configure policing mappings, enter:

   ```
   [no] policing-cos-map ip <dscp-value> <conformance-level>

   [drop-precedence <drop-precedencevalue>]

   {[user-priority <user-priority-value>] | [dscp <dscp-value>]}
   ```

Table 76: Variable definitions

Variable	Value
<dscp-value>	The DSCP value in the IP packet. Valid Range: 0- 63
<conformance-level>	The conformance-level assigned by the policer to the packet. Valid Range: conform / exceed / violate
[drop-precedence <drop-precedencevalue>]	Optional. The drop precedence value to be assigned. Valid Range: low / medium / high
[user-priority <user-priority-value>]	Optional. The user priority value to be assigned. Valid Range: 0-7

Variable	Value
[dscp <dscp-value>]	Optional. The DSCP value to be assigned. Valid Range: 0-63
[no]	resets the CoS mappings for the packet

Disabling policing for the class

By default, after you configure srTCM or trTCM properties for a class, the configuration is automatically enabled for the class. You can disable policing for a class.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To specify the policy map, enter:

   ```
   policy-map <policy-map-name>
   ```

4. To specify the class map, enter:

   ```
   class-map <classmap-name>
   ```

5. To specify police configuration, enter:

   ```
   police
   ```

6. To disable policing for the class, enter:

   ```
   [no] disable
   ```

Table 77: Variable definitions

Variable	Value
[no]	Enables policing for the class

Configuring Single Rate Three Color Marker

Configure srTCM to configure single-rate policing for the class. By default, after you configure srTCM properties, the configuration is automatically enabled for the class.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Module QoS, enter:

```
qos module
```

3. To specify the policy map, enter:

```
policy-map <policy-map-name>
```

4. To specify the class map, enter:

```
class-map <classmap-name>
```

5. To specify police configuration, enter:

```
police
```

6. To configure srTCM, enter:

```
srtcm <cir-value> [cbs <cbs-value>] [ebs <ebs-value>]
```

Table 78: Variable definitions

Variable	Value
<cir-value>	The committed information rate. Valid Range: 3Kbps - 10000000 Kbps
[cbs <cbs-value>]	Optional. The committed burst size in Kbits. The valid range corresponding to CIR should be >= (1 * CIR Kbps) and <= (5* CIR Kbps). The value must be at least as large as the largest possible packet.
[ebs <ebs-value>]	Optional. The excess burst size in Kbits. The valid range corresponding to CIR should be >= (1 * CIR Kbps) and <= (5* CIR Kbps). The value must be at least as large as the largest possible packet.
[no]	Removes srTCM configuration for the class

Configuring Two Rate Three Color Marker

Configure trTCM to configure two-rate policing for the class. By default, after you configure trTCM properties, the configuration is automatically enabled for the class.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Module QoS, enter:

```
qos module
```

3. To specify the policy map, enter:

```
policy-map <policy-map-name>
```

4. To specify the class map, enter:

```
class-map <classmap-name>
```

5. To specify police configuration, enter:

```
police
```

6. To configure trTCM, enter:

```
trtcm <cir-value> <pir-value> [cbs <cbs-value>] [pbs <pbs-value>]
```

Table 79: Variable definitions

Variable	Value
<cir-value>	The committed information rate. Valid Range: 3 Kbps - 10000000 Kbps.
<pir-value>	The peak information rate. Valid Range: 3 Kbps - 10000000 Kbps.
[cbs <cbs-value>]	Optional. The committed burst size. The valid range corresponding to CIR should be >= (1 * CIR Kbps) and <= (5* CIR Kbps). The value must be at least as large as the largest possible packet.
[pbs <pbs-value>]	Optional. The peak burst size. The valid range corresponding to CIR should be >= (1 * CIR Kbps) and <= (5* CIR Kbps). The value must be at least as large as the largest possible packet.
[no]	Removes the trTCM configuration for the class.

Configuring color aware mode for srTCM and trTCM

Enable srTCM or trTCM to run in color aware mode for the class map to allow consideration of previous DP markings.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To specify Module QoS, enter:

```
qos module
```

3. To specify the policy map, enter:

```
policy-map <policy-map-name>
```

4. To specify the class map, enter:

   ```
   class-map <classmap-name>
   ```

5. To specify police configuration, enter:

   ```
   police
   ```

6. To configure color-aware mode, enter:

   ```
   [no] color-aware
   ```

Table 80: Variable definitions

Variable	Value
[no]	Configures color blind operation for the policer

Configuring CoS re-marking for the policer assigned to the class

Enable CoS re-marking for the policer assigned to the class

By default, Policing based CoS re-marking is disabled.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To specify the policy map, enter:

   ```
   policy-map <policy-map-name>
   ```

4. To specify the class map, enter:

   ```
   class-map <classmap-name>
   ```

5. To specify police configuration, enter:

   ```
   police
   ```

6. To enable or disable CoS remarking, enter:

   ```
   [no] remark-cos
   ```

Table 81: Variable definitions

Variable	Value
[no]	disable policing based CoS remarking for the class

Configuring packet drop for violating packets

Enable packet drop for packets that violate the policing profile

By default, packet drop is disabled.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```
2. To specify Module QoS, enter:

   ```
   qos module
   ```
3. To specify the policy map, enter:

   ```
   policy-map <policy-map-name>
   ```
4. To specify the class map, enter:

   ```
   class-map <classmap-name>
   ```
5. To specify police configuration, enter:

   ```
   police
   ```
6. To configure packet drop, enter:

   ```
   [no] drop-violate
   ```

Table 82: Variable definitions

Variable	Value
[no]	Removes dropping of policing profile violated packets

Enabling Accounting

Enable QoS accounting to obtain flow-based accounting of the number of forwarded, conforming, exceeding, and violated packets for the class.

By default, accounting is disabled for the class

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```
2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To specify the policy map, enter:

   ```
   policy-map <policy-map-name>
   ```

4. To specify the class map, enter:

   ```
   class-map <classmap-name>
   ```

5. To enable or disable accounting for the class, enter:

   ```
   [no] accounting enable
   ```

Table 83: Variable definitions

Variable	Value
[no]	Disables accounting for the class.

Disabling accounting

Disable QoS accounting for billing counters at the system level.

By default, accounting is enabled at the system level.

Important:

The SR2330 does not support global enable and disable of the accounting feature. Accounting and billing can be enabled and disabled for each class map, similar to the SR4134.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To configure accounting, enter:

   ```
   [no] accounting disable
   ```

Table 84: Variable definitions

Variable	Value
[no]	Enables accounting at the system level

Configuring flow rate monitoring

1. Configure flow rate monitoring for classes.
2. Configure the global sampling interval and sampling period.
3. Enable rate sampling at the system level to initiate the flow rate monitoring of all classes configured for monitoring.

Configuring flow rate monitoring for a class

Configure flow rate monitoring for a class to monitor the class traffic flows. By default, rate monitoring is disabled for a new class.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To specify the policy map, enter:

   ```
   policy-map <policy-map-name>
   ```

4. To specify the class map, enter:

   ```
   class-map <classmap-name>
   ```

5. To configure flow rate monitoring for the class, enter:

   ```
   [no] rate-monitoring enable
   ```

Table 85: Variable definitions

Variable	Value
[no]	Disables rate-monitoring for the class.

Configuring global rate monitoring parameters

Configure rate-monitoring parameters at the global level to specify the sampling interval and sampling period.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To configure global rate monitoring parameters, enter:

   ```
   [no] rate-monitoring [sampling-interval <interval>]
   [sampling-period <period>]
   ```

Table 86: Variable definitions

Variable	Value
[sampling-interval <interval>]	Optional. Sampling interval in seconds. Valid range: 1 – 300
[sampling-period <period>]	Optional. Sampling period in seconds. Valid range: 1 – 4294967295
[no]	Resets the flow-rate monitoring parameters to default

Enabling rate sampling at the system level

Enable rate sampling at the system level to initiate the flow rate monitoring of all classes configured for monitoring.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To configure rate sampling, enter:

   ```
   [no] enable-rate-sampling
   ```

Table 87: Variable definitions

Variable	Value
[no]	Disables rate sampling.

Configuring Congestion management

Configuring congestion avoidance parameters

Configure the congestion avoidance parameters for one of the four congestion profiles.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To specify the congestion profile to configure, enter:

   ```
   congestion-profile <1-4>
   ```

4. To configure congestion profile parameters, enter:

   ```
   [no] drop-curve <queue-id> <drop-precedence-value>
   [maxthreshold <max-threshold-value>] [min-threshold
   <minthreshold- value>] [drop-prob-denominator <mpd>]
   ```

Table 88: Variable definitions

Variable	Value
<1-4>	Specifies the congestion profile.
<queue-id>	Identifier for the queue. Valid ranges: 1-8
<drop-precedence-value>	Drop precedence value. Valid ranges: {low, min, high}
[maxthreshold <max-threshold-value>]	Optional. Maximum threshold value for WRED. Valid range: 1 - 4000
[min-threshold <minthreshold- value>]	Optional. Minimum threshold value for WRED. Valid range: 1 - 4000
[drop-prob-denominator <mpd>]	Optional. Mark probability denominator. Valid range: 1 - 10
[no]	Disables accounting for the class.

Configuring EWF for profile

Configure the exponential weighting factor (EWF) constant for the specified congestion profile.

Important:

The SR2330 only supports two congestion management schemes: Simple Random Early Drop (SRED) and Tail drop. Unlike the SR4134, WRED is not supported. As a result, the command **`exponential-weighting-constant`** is not applicable and not supported.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify Module QoS, enter:

   ```
   qos module
   ```

3. To specify the congestion profile to configure, enter:

   ```
   congestion-profile <1-4>
   ```

4. To configure the exponential weighting factor, enter:

   ```
   [no] exponential-weighting-constant <queue-id> <ewf-value>
   ```

Table 89: Variable definitions

Variable	Value
<queue-id>	Identifier for the queue. Valid ranges: 1-8
<ewf-value>	EWF value. Valid ranges: 0 – 10
[no]	Resets the EWF parameter to default value.

Assigning congestion profile to an interface

Attach the congestion profile to an interface to apply congestion management to the interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

```
interface ethernet <slot/port>
```

3. To specify Module QoS, enter:

```
qos module
```

4. To clone the congestion profile to the interface, enter:

```
[no] congestion-profile <profile-no>
```

Table 90: Variable definitions

Variable	Value
<profile-no>	The congestion profile to apply. Valid ranges: 1-4
[no]	Removes the profile from the interface.

Enabling and disabling RED

Disable RED for an interface to enable tail drop on the interface. If tail drop is no longer required, you can re-enable RED on the interface.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select the interface, enter:

```
interface ethernet <slot/port>
```

3. To specify Module QoS, enter:

```
qos module
```

4. To enable or disable RED, enter:

```
[no] random-detect
```

Table 91: Variable definitions

Variable	Value
[no]	Disables random-early-detect on the interface

Configuring Queueing and Scheduling

Configuring Strict Priority scheduling

Configure the queue for Strict Priority scheduling to enable interrupt queuing.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To select the queue to configure, enter:

   ```
   queue <1-8>
   ```

5. To set the selected queue to strict priority queuing, enter:

   ```
   priority-queue
   ```

Configuring Weighted Round Robin scheduling

Configure the queue as part of Weighted Round Robin scheduling group to provide WRR scheduling for the queue group.

Important:

The SR2330 supports only one Weighted Round Robin (WRR) group for the interface queues unlike the SR4134 which supports two WRR groups. As a result, on the SR2330 the **`wrr-group`** option in the **`wrr-queue`** command is ignored.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To select the queue to configure, enter:

   ```
   queue <1-8>
   ```

5. To configure the selected queue for DWRR, enter:

   ```
   wrr-queue <queue-weight> <WRR-group>
   ```

Table 92: Variable definitions

Variable	Value
<queue-weight>	Specifies the relative queue weight (1-255).
<WRR-group>	Specifies the WRR group to which the queue belongs. Within the same WRR group, queues must be consecutive.

Configuring maximum queue limit

Configure the maximum queue limit for a queue:

Important:

The SR2330 does not support excess buffer-limit configuration (using the **`queue-limit`** command) for an interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To select the queue to configure, enter:

   ```
   queue <1-8>
   ```

5. To configure the queue limit, enter:

   ```
   queue-limit <16-4000>
   ```

Table 93: Variable definitions

Variable	Value
<16-4000>	Queue limit in packets. Specify in multiples of 16.

Configuring traffic shaping

Configuring queue-based shaping

Configure the shaping rate and burst size for the queue to configure queue-based shaping.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To select the queue to configure, enter:

   ```
   queue <1-8>
   ```

5. To configure shaping for the queue, enter:

   ```
   [no] shape [rate <rate>] [burst <burst>]
   ```

Table 94: Variable definitions

Variable	Value
[rate <rate>]	Queue shaping rate in Kb/s. Valid range: 651 - 999936 in multiples of 651 Kb/s.
[burst <burst>]	Burst size in bytes in multiples of 4KB. Valid range: 4 – 16000
[no]	disables shaping

Configuring port-based shaping

Configure the shaping rate and burst size for the port.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure shaping for the port, enter:

   ```
   [no] shape [rate <rate>] [burst <burst>]
   ```

Table 95: Variable definitions

Variable	Value
[rate <rate>]	Queue shaping rate in Kb/s.Valid range: 651 - 999936 in multiples of 651 Kb/s.
[burst <burst>]	Burst size in bytes in multiples of 4KB. Valid range: 4 – 16000
[no]	disables shaping

Configuring Buffer Management

Configuring receive buffers on an ingress port

Configure the receive buffers on an ingress port

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure buffers, enter:

   ```
   ingress-buffer-limit <buffers>
   ```

Table 96: Variable definitions

Variable	Value
<buffers>	Buffer count. Valid Range: 0 – 4000

Configuring transmit descriptors on an egress port

Configure the transmit descriptors on an egress port.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure transmit descriptors, enter:

   ```
   egress-buffer-limit <buffers>
   ```

Table 97: Variable definitions

Variable	Value
<buffers>	Buffer count. Valid Range: 0 – 4000

Configuring XOFF threshold limit on an interface

Configure the XOFF threshold limit for an interface

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure XOFF threshold limit, enter:

```
xoff-limit <buffers>
```

Table 98: Variable definitions

Variable	Value
<buffers>	Buffer count. Valid Range: 0 – 4000

Configuring XON threshold limit on an interface

Configure the XON threshold limit for an interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify Module QoS, enter:

   ```
   qos module
   ```

4. To configure XON threshold limit, enter:

   ```
   xon-limit <buffers>
   ```

Table 99: Variable definitions

Variable	Value
<buffers>	Buffer count. Valid Range: 0 – 4000

Disabling QoS globally

By default, QoS is enabled in the system.

Enable or disable QoS for all the applicable chassis and module interfaces.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify QoS configuration, enter:

```
qos
```

3. To configure QoS global status, enter:

```
[no] disable-qos
```

Table 100: Variable definitions

Variable	Value
[no]	Enables QoS at global level.

Disabling QoS on an interface

By default, QoS is enabled on the interface

Enable or disable QoS for the specified interface.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select the interface, enter:

```
interface ethernet <slot/port>
```

3. To specify QoS configuration, enter:

```
qos
```

4. To configure QoS status on the interface, enter:

```
[no] disable-qos
```

Table 101: Variable definitions

Variable	Value
[no]	Enables QoS at interface level.

Configuring auto QoS

Enabling auto QoS globally

Enable auto QoS at the global level to apply auto QoS to all the applicable interfaces (chassis and module interfaces). If an interface already has a policy map applied, auto QoS is not applied to that interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To specify QoS configuration, enter:

   ```
   qos
   ```

3. To configure auto QoS globally for Ethernet modules, enter:

   ```
   [no] enable-auto-qos
   ```

Table 102: Variable definitions

Variable	Value
[no]	Disables auto qos at the system level

Enabling auto QoS on the interface

Enables auto QoS at interface level

By default, auto-qos will be disabled on any interface.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the interface, enter:

   ```
   interface ethernet <slot/port>
   ```

3. To specify QoS configuration, enter:

```
qos
```

4. To configure auto QoS for the interface, enter:

```
[no] enable-auto-qos
```

Table 103: Variable definitions

Variable	Value
[no]	Disables auto qos at the interface level

Displaying Module QoS configuration and statistics

Displaying the system-level Module QoS configuration

Display the system level Module QoS configuration

Procedure steps

To display the system QoS configuration, enter:

```
show qos module system
```

Displaying the interface-level Module QoS configuration

Display the interface-level Module QoS configuration.

Procedure steps

To display the interface QoS configuration, enter:

```
show qos module ethernet <slot/port>
```

Table 104: Variable definitions

Variable	Value
<slot/port>	Slot/port of the interface

Displaying module QoS policy map and class map configuration

Display the module QoS policy map and class map configuration.

Procedure steps

To display policy map and class map, enter:

```
show qos module policy-map [<policy-name>] [class <class-name>]
[detail]
```

Table 105: Variable definitions

Variable	Value
[<policy-name>]	Optional. Name of the policy-map. Valid range: Up to 19 characters. If policy-name is not specified, the QoS configuration of all the policy maps in the system will be displayed.
[class <class-name>]	Optional. Name of the class-map. Valid range: Up to 19 characters If class-name is not specified, the QoS configuration of all the class maps within the policy map will be displayed
[detail]	Optional. Used for detailed display of the configuration User may opt to display the detailed information using the “detail” parameter

Displaying the interface policy map

Display the policy map attached with the interface.

Procedure steps

To display the interface policy map, enter:

```
show qos module service-policy [<if-name>]
```

Table 106: Variable definitions

Variable	Value
[<if-name>]	Optional. Name of the interface. If if-name is not specified, all the policy map and interface mappings in the system will be displayed.

Displaying the non-IP policy CoS mappings

Display the policy CoS mappings configured for the non-IP packets.

Procedure steps

To display the non-ip policy CoS mappings, enter:

```
show qos module policing-cos-map non-ip [<user-priority-
value>]
```

Table 107: Variable definitions

Variable	Value
[<user-priority- value>]	Optional. The user priority value in the priority tagged packet. Valid Range: 0-7 If user-priority-value is specified, all the policing-cos-mappings associated with the user-priority value will be displayed. Otherwise, displays all the mappings

Displaying the IPv4 policy CoS mappings

Display the policy CoS mappings configured for the IP packets.

Procedure steps

To display the IPv4 policy CoS mappings, enter:

```
show qos module policing-cos-map ip [<dscp-value>]
[<conformance-level>]
```

Table 108: Variable definitions

Variable	Value
[<dscp-value>]	Optional. The DSCP value in the IP packet. Valid Range: 0-63 If dscp-value is specified, all the policing-cos-mappings associated with the DSCP value will be displayed. Otherwise, displays all the mappings
[<conformance-level>]	Optional. The conformance-level assigned by the policer to the packet. Valid Range: conform / exceed / violate. If conformance-level is specified, the policing-cos-mapping associated with the DSCP value and the conformance level specified will be displayed. Otherwise, all the policing-cos-mappings associated with the DSCP value specified will be displayed.

Displaying congestion profiles

Display the congestion avoidance parameters configured for the congestion profile.

Procedure steps

To display congestion profile, enter:

```
show qos module congestion-profile <profile-id> [<queue-id>]
```

Table 109: Variable definitions

Variable	Value
<profile-id>	One of the four congestion profile IDs. Valid range: 1 – 4
[<queue-id>]	Optional. Queue identifier. Valid range: 1 – 8 If queue-id is not specified, the congestion avoidance parameters associated with all 8 queues in the congestion profile will be displayed.

Displaying system-level QoS status

Display the status of system-level QoS configuration (QoS enabled/disabled and auto-QoS enabled/disabled).

Procedure steps

To display system-level QoS configuration, enter:

```
show qos system
```

Clearing QoS counters

Clear the counters associated with the specified policy map and classes.

Procedure steps

To clear QoS counters, enter:

```
clear qos module policy-map <policy-name> [class <class-name>]
[rate-monitoring-stats]
```

Table 110: Variable definitions

Variable	Value
<policy-name>	Name of the policy-map.
[class <class-name>]	Optional. Name of the class-map If class-name is not specified, the counters associated with all the class maps within the policy-map specified will be cleared
[rate-monitoring-stats]	Optional. Specifies the rate monitoring statistics If rate-monitoring stats is specified, the rate-monitoring stats associated with the class map will be cleared

Chapter 9: SLA configuration

Refer to the following procedures to configure SLA.

Be aware that the total number of SLA operations configured for the same destination is restricted to 5. The 5 entries can be of any SLA type for the same destination (even same types can be repeated). This restriction is applicable for both IPv4 and IPv6 destinations.

Creating an SLA profile

Create an SLA profile.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To configure the SLA profile, enter:

   ```
   [no] sla profile <profile-id>
   ```

Table 111: Variable definitions

Variable	Value
<profile-id>	ID of the profile to be configured. Valid range is 1- 1000.
[no]	Removes the SLA profile.

Configuring a schedule for an SLA profile

Configure a schedule for the SLA profile.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

```
sla profile <profile-id>
```

3. To configure the SLA profile schedule, enter:

```
[no] sla schedule <sla-profile> [life <time-inmin>] [start
<time-in-min> ]
```

Table 112: Variable definitions

Variable	Value
<sla-profile>	This parameter specifies which sla profile has to be scheduled. Valid : 1-1000
[life <time-inmin>]	It is the time until which the sla profile will be in Active state. Valid: 1 – 1440. Default value: Forever (means SLA will be active forever till no sla schedule command is given)..
[start <time-in-min>]	This specifies the time after which the profile will be scheduled. Valid: 0 – 60. Default value: 0.
[no]	Stop the scheduling of the SLA profile.

Specifying a description for the SLA profile

Configure description for the SLA profile.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select the SLA profile, enter:

```
sla profile <profile-id>
```

3. To configure the description for the SLA profile, enter:

```
[no] description <name>
```

Table 113: Variable definitions

Variable	Value
<name>	This parameter specifies the description for the SLA profile.
[no]	Removes the description.

Configuring UDP echo

To calculate RTT for UDP packets to the destination system, configure the IPv4 address and port number of the destination system.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure the UDP echo, enter:

   ```
   [no] udp-echo <ip-address> [port <7|50001>]
   ```

Table 114: Variable definitions

Variable	Value
<ip-address>	IP address of the destination system.
[port <7|50001>]	The port to which the SLA packets will be sent. Valid: 7 or 50001. Default: 50001
[no]	Removes all the udp-echo configurations.

Configuring UDPv6 echo

To calculate RTT delay for UDPv6 packets to the destination system, configure the IPv6 address and port number of destination system.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure the UDPv6 echo, enter:

   ```
   [no] udp-v6-echo <ipv6-address> [port <7|50002>]
   ```

Table 115: Variable definitions

Variable	Value
<ipv6-address>	IPv6 address of the destination system.
[port <7|50002>]	The port to which the SLA packets will be sent. Valid: 7 or 50002. Default: 50002
[no]	Removes all the udp-v6-echo configurations.

Configuring ICMP echo

To calculate RTT delay, configure ICMP echo parameters.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure the ICMP echo, enter:

   ```
   [no] icmp-echo <ip-address>
   ```

Table 116: Variable definitions

Variable	Value
<ip-address>	IP address of the destination system.
[no]	Removes all the icmp-echo configurations.

Configuring ICMPv6 echo

Configure ICMPv6 echo.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

```
sla profile <profile-id>
```

3. To configure ICMPv6 echo, enter:

```
[no] icmp-v6-echo <ipv6-address>
```

Table 117: Variable definitions

Variable	Value
<ipv6-address>	IPv6 address of the destination system.
[no]	Removes all the icmp-v6-echo configurations.

Configuring UDP jitter

To calculate two-way/one-way delay/jitter, configure the IPv4 address of the destination and other related parameters.

Procedure steps

1. To enter the configuration mode, enter:

```
configure terminal
```

2. To select the SLA profile, enter:

```
sla profile <profile-id>
```

3. To configure UDP jitter, enter:

```
[no] udp-jitter <operation> <ip-addr>
[port <port>]
[packet-spacing <value>]
[packet-size <value>]
[packet-count <value>]
[periodicity <value>]
{[dscp <value>] | [ip-precedence <value>] | [tos-byte
<value>]}
[ttl <value>]
```

Table 118: Variable definitions

Variable	Value
<operation>	The type of parameter it has to measure and monitor. Valid: delay, jitter and packet-loss

Variable	Value
<ip-addr>	IP address of the destination system.
[port <port>]	The port to which the packets have to be sent. Valid: 7 or 50001. Default: 50001.
[packet-spacing <value>]	Time spacing between two consecutive packets in mill seconds. Valid: 10 – 5000. Default: 20
[packet-size <value>]	Size of the packet in bytes. Valid: 1 – 1400. Default: 80
[packet-count <value>]	Number of packets to be sent. Valid: 1 – 100. Default: 20
[periodicity <value>]	After every specified periodicity minutes the sla will be scheduled again. Valid: 10 – 60. Default: 10
[dscp <value>]	DSCP value to be marked on the packets. Valid: 0 – 63. Default: 0
[ip-precedence <value>	IP-precedence to be marked. Valid: 0 – 7. Default: 0
[tos-byte <value>]	ToS for the packets. Valid: 0 – 255. Default: 0
[ttl <value>]	Time to Live for the packet. Valid: 1 – 255. Default: 64
[no]	Removes all the udp-jitter configurations.

Configuring UDPv6 jitter

To calculate two-way/one-way delay/jitter for UDPv6 packets, configure the IPv6 address of destination and other related parameters.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure UDP jitter, enter:

   ```
   [no] udp-v6-jitter <operation> <ipv6-addr>
   [port <port>]
   [packet-spacing <value>]
   [packet-size <value>]
   [packet-count <value>]
   ```

```
[periodicity <value>]
[tc <value>]
[hop-limit <value>]
[flow-label <value>]
```

Table 119: Variable definitions

Variable	Value
<operation>	The type of parameter it has to measure and monitor. Valid: delay, jitter and packet-loss
<ipv6-addr>	IPv6 address of the destination system.
[port <port>]	The port to which the packets have to be sent. Valid: 7 or 50002. Default: 50002.
[packet-spacing <value>]	Time spacing between two consecutive packets in mill seconds. Valid: 10 – 5000. Default: 20
[packet-size <value>]	Size of the packet in bytes. Valid: 1 – 1400. Default: 80
[packet-count <value>]	Number of packets to be sent. Valid: 1 – 100. Default: 20
[periodicity <value>]	After every specified periodicity minutes the sla will be scheduled again. Valid: 10 – 60. Default: 10
[tc <value>]	Traffic class of the IPv6 header. Valid: 0 – 255. Default: 0
[hop-limit <value>]	Hop limit for the IPv6 packet. Valid: 1 – 255. Default: 64
[flow-label <value>]	Flow label for the packet Valid: 0 – 1048575. Default: 0
[no]	Removes all the udp-v6-jitter configurations.

Configuring actions for SLA event handling

Configure SLA event handling to configure the effect type and action type for the SLA.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure actions, enter:

   ```
   [no] action <effect-type> [<action-type>]
   ```

Table 120: Variable definitions

Variable	Value
<effect-type>	The SLA variable that has to be monitored. Valid: jitter-average,jitter-average-src-dest,jitter-average-dest-src,jittermax- positive-src-dest,jitter-max-positive-dest-src,jitter-maxnegative- src-dest,jitter-max-negative-dest-src,delayaverage, delay-average-src-dest,delay-average-dest-src,delaymax- src-dest,delay-max-dest-src,packet-loss,packet-out-oforder, packet-late-arrival, response-time or response-time-average.
[<action-type>]	The action to be taken when the monitored variable exceeds the range. Valid: console-logging (default), syslog or trap.
[no]	Makes the effect type as null.

Configuring the threshold violation type

Configure threshold violation type.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure the threshold type, enter:

   ```
   [no] threshold-type <type>
   ```

Table 121: Variable definitions

Variable	Value
<type>	Threshold type to be configured. Valid: immediate, average, consecutive or xofy.
[no]	Makes the threshold type as null.

Configuring the threshold value

Configure threshold values.

Procedure steps

1. To enter the configuration mode, enter:

   ```
   configure terminal
   ```

2. To select the SLA profile, enter:

   ```
   sla profile <profile-id>
   ```

3. To configure the threshold values, enter:

   ```
   [no] threshold-value <value1> <value2>
   ```

Table 122: Variable definitions

Variable	Value
<value1>	Threshold Value1. Valid: 1-10000
<value2>	Threshold Value2. Valid: 1-10000
[no]	Sets the threshold values to null.

Displaying the SLA profile

Display the configured SLA profile.

Procedure steps

To display the configured SLA profile, enter:

```
show sla profile [<1-1000>] [detail]
```

Table 123: Variable definitions

Variable	Value
[<1-1000>]	Specifies the SLA profile to display.
[detail]	Displays detailed information about the SLA profile.

Clearing the SLA statistics

Clear the statistics for the SLA profile.

Procedure steps

To clear the configured SLA profile, enter:

```
clear sla profile [<1-1000> | all]
```

Table 124: Variable definitions

Variable	Value
[<1-1000> | all]	Specifies the SLA profile statistics to clear.

Chapter 10: Configuration examples

Chassis QoS configuration examples

Multifield classification example

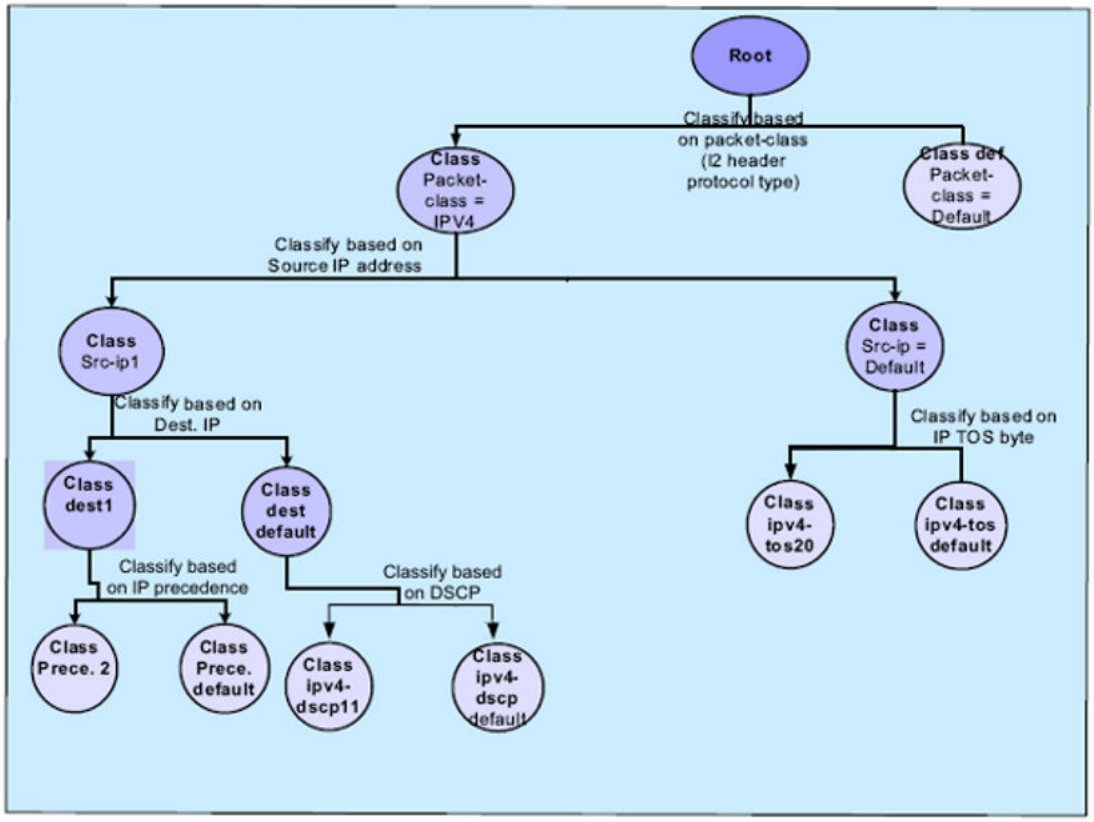

Figure 14: Multifield classification example

The following commands describe how to configure the multifield classification example shown in the preceding figure.

Procedure steps

1. To specify chassis QoS configuration, enter:

```
configure terminal
qos chassis
```

2. To specify the policy map to configure, enter:

```
policy-map xyz
```

3. To configure the IPv4 packet class (packet-class), enter:

```
class-map packet-class root match packetclass ipv4
cbq pr-percent 100
exit
```

4. To configure the default protocol type packet class (packet-default), enter:

```
class-map packet-default root
match packetclass default
cbq cr-percent 5 pr-percent 100 priority 8
police
srtcm cir-percent 10
exit
exit
```

5. To configure the source IP 1 packet class (src-ip1), enter:

```
class-map src-ip1 packet-class
match src-address 1.2.3.4
match src-address 1.2.5.0/24
match src-address 6.4.7.8-6.4.7.200
cbq pr-percent 100
exit
```

6. To configure the default source IP packet class (src-ip), enter:

```
class-map src-ip packet-class
match src-address default
cbq pr-percent 100
exit
```

7. To configure the destination IP packet class (dest1), enter:

```
class-map dest1 src-ip1
match dest-address 20.1.1.1
match dest-address 20.1.2.0/24
match dest-address 20.1.3.10-20.1.3.50
cbq pr-percent 100
exit
```

8. To configure the default destination IP packet class (dest), enter:

```
class-map dest src-ip1
match dest-address default
cbq pr-percent 100
exit
```

9. To configure the precedence packet class (prece2), enter:

```
class-map prece2 dest1
match precedence 2
cbq cr-percent 15 prpercent 100 priority 2
police
trtcm cirpercent 20 pir-percent 40
```

```
exit
exit
```

10. To configure the default precedence packet class (prece), enter:

```
class-map prece dest1
match precedence default
cbq cr-percent 10 prpercent 100 priority 7
police
srtcm cir-percent 10
exit
exit
```

11. To configure the DSCP packet class (dscp11), enter:

```
class-map ipv4-dscp11 dest
match dscp af11
cbq cr-percent 10 pr-percent 100 priority 6
police
trtcm cirpercent 10 pir-percent 50
exit
exit
```

12. To configure the default DSCP packet class (dscp), enter:

```
class-map ipv4-dscp dest
match dscp default
cbq cr-percent 10 pr-percent 100 priority 7
police
srtcm cirpercent 10
exit
exit
```

13. To configure the ToS packet class, enter:

```
class-map ipv4-tos20 src-ip
match tos 20
cbq cr-percent 10 pr-percent 100 priority 2
police
srtcm cirpercent 10
exit
exit
```

14. To configure the default ToS packet class, enter:

```
class-map ipv4-tos src-ip
match tos default
cbq cr-percent 10 prpercent 100 priority 7
police
srtcm cirpercent 10
exit
exit
pop
```

15. To apply the policy map to an interface, enter:

```
interface bundle wan1
qos chassis
service-policy output xyz
```

Class of Service Marking

After you have classified your traffic using class maps, you can use Class of Service marking commands to mark the traffic.

In the following example, the commands shown are used to remark incoming CoS values to DSCP values.

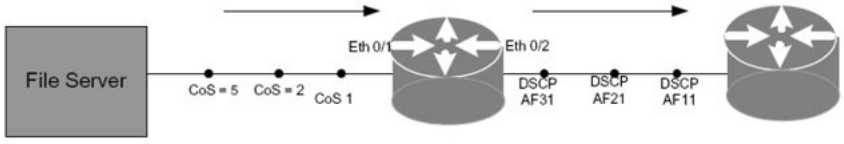

Figure 15: CoS to DSCP remarking

Procedure steps

1. To specify the policy map to configure, enter:

```
configure terminal
qos chassis
policy-map <policy-name>
```

2. To configure the CoS marking, enter:

```
class-map voice
match tos 5
mark dscp af31
exit

class-map data
match tos 2
mark dscp af21
SR4134/configure/qos/chassis/policy-map <policy-name>/class-map data#
exit

class-map default
match tos 1
mark dscp af11
exit
```

In this example, traffic marked with ToS values of 5 is classified in the voice class map, whereas traffic with ToS values of 2 goes into the data class map. ToS values of 1 are placed in the default class map. The DSCP marking assigns a DSCP value of AF31 to the voice traffic, a DSCP value of AF21 to the data traffic, and a DSCP value of AF11 to the default traffic.

Congestion management

In the following congestion management example, the goal is to configure WRED for all three classes voice, data and default created in policy-map profile-out. In this example the behavior is consider only for the green packets. Each class has its own egress queues.

Procedure steps

1. To specify the policy map to configure, enter:

```
configure terminal
qos chassis
policy-map profile-out
```

2. To configure congestion management, enter:

```
class-map default
red minthgreen 25 maxth-green 100 mpd-green 4
exit

class-map voice
red minthgreen 40 maxth-green 100 mpd-green 4
exit

class-map data
red minth-green 50 maxth-green 100 mpd-green 4
exit
```

The default traffic starts to discard the packets only when the queue size reaches 25%. Voice traffic is not discarded until the queue size reaches 40% of the queue size, and finally, data traffic is not discarded until the queue size reaches 50%. If the queue size exceeds 100%, there is a 100 percent chance of discard for these three types of traffic. However, when the queue size is exactly 100%, the chance of discard for these various packet types is 25 percent.

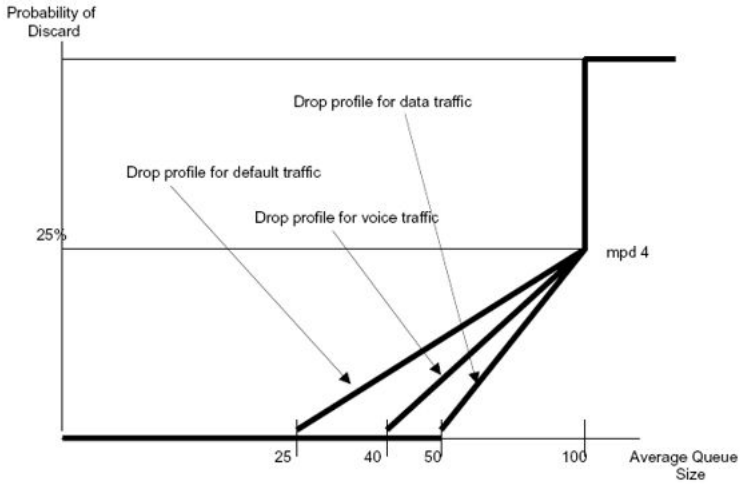

Figure 16: Discard probability in relation to average queue size

The MPD value of 4 is chosen to meet the requirement of a 25 percent chance discard when the queue size equals the maximum threshold (that is, 1/4 = .25). Also, notice that default traffic

is dropped before voice traffic, which is dropped before data. This approach is consistent with the definition of the per-hop behaviors (PHBs).

MPLS QoS

The following is an example of configuration of global DSCP-to-EXP mapping table.

Procedure steps

1. To specify Chassis QoS configuration, enter:

```
configure terminal
qos chassis
```

2. To configure global DSCP-to-EXP mapping, enter:

```
dscp-exp-cos-map cs7 7
dscp-exp-cos-map cs6 6
dscp-exp-cos-map cs5 5
dscp-exp-cos-map ef 5
dscp-exp-cos-map cs4 4
dscp-exp-cos-map af41 4
dscp-exp-cos-map af42 4
dscp-exp-cos-map af43 4
dscp-exp-cos-map cs3 3
dscp-exp-cos-map af31 3
dscp-exp-cos-map af32 3
dscp-exp-cos-map af33 3
dscp-exp-cos-map cs2 2
dscp-exp-cos-map af21 2
dscp-exp-cos-map af22 2
dscp-exp-cos-map af23 2
dscp-exp-cos-map cs1 1
dscp-exp-cos-map af11 1
dscp-exp-cos-map af12 1
dscp-exp-cos-map af13 1
dscp-exp-cos-map cs0 0
```

The following is the example of configuration of flow based EXP marking. Note that the flow based EXP marking is allowed in the inbound direction only.

Configure the policy map with required flows and associated EXP markings, then apply the policy map to an interface. .

Procedure steps

1. To specify Chassis QoS configuration, enter:

```
configure terminal
qos chassis
```

2. To create the policy map, enter:

```
policy-map xyz
```

3. To configure the packet-class class map, enter:

```
class-map packet-class root
 match packetclass ipv4
```

```
cbq pr-percent 100
exit
```

4. To configure the packet-default class map with EXP marking, enter:

```
class-map packet-default root
match packetclass default
cbq cr-percent 5 pr-percent 100 priority 8
exit
```

5. To configure the src-ip1 leaf class map with EXP marking, enter:

```
class-map src-ip1 packet-class
match src-address 1.2.3.4
match src-address 1.2.5.0/24
match src-address 6.4.7.8-6.4.7.200
cbq pr-percent 100
mark exp 4
exit
```

6. To configure the src-ip2 leaf class map with EXP marking, enter:

```
class-map src-ip2 packet-class
match src-address 8.7.3.4
match src-address 8.7.5.0/24
match src-address 11.6.7.8-11.6.7.200
cbq pr-percent 100
mark exp 3
exit
```

7. To configure the src-ip leaf class map with EXP marking, enter:

```
class-map src-ip packet-class
match src-address default
cbq pr-percent 100
mark exp 2
exit
pop
```

8. To apply the policy map on the inbound direction of an interface, enter:

```
interface ethernet 0/2
qos chassis
service-policy input xyz
```

Module QoS configuration example

The following figure shows a typical configuration scenario.

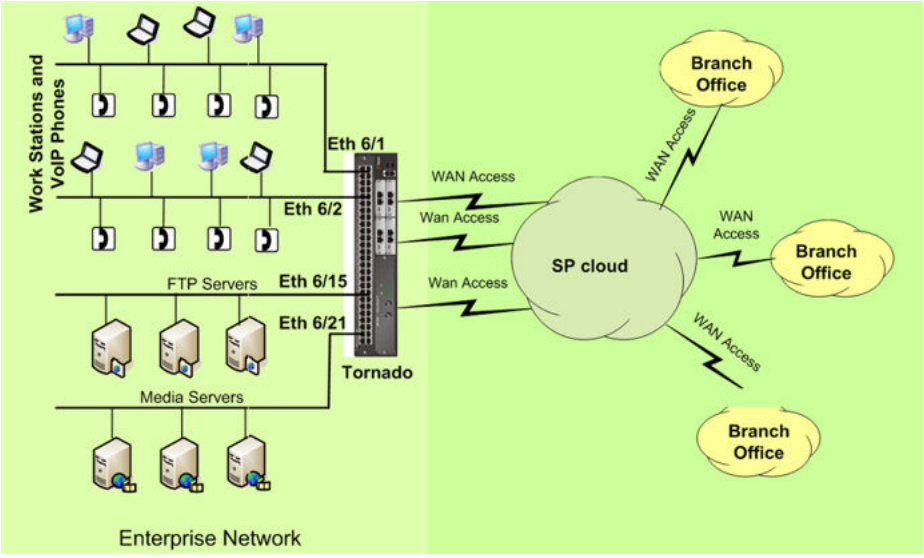

Figure 17: Module QoS configuration example

The following sections describe an example QoS configuration for this scenario.

Configuration for interactive voice

Procedure steps

1. Specify Module QoS configuration

```
configure terminal
qos module
```

2. Configure a policy map and class for voice

```
policy-map enterprise
class-map voice
```

3. Configure rules to classify RTP packets (assuming port 50000 is used for voice)

```
match ipv4 protocol 17 src-port 50000-55000
match ipv4 protocol 17 dest-port 50000-55000
```

4. Mark voice traffic with EF DSCP so that other succeeding devices provide the right QoS

```
mark-dscp 46
```

5. Police voice traffic to (90Kbps * 10 users) and drop violating packets. Shaping is not suitable for voice traffic

```
police
srtcm 900
drop-violate
exit
```

6. Assign a strict priority queue with high priority to the voice traffic to minimize delay and jitter (by default the packet processor is configured with 3 priority queues 1, 2 and 3 in decreasing order of priority)

```
assign-queue 1
exit
pop
```

7. Apply the policy map to applicable interfaces

```
interface ethernet 6/1
qos module
service-policy input enterprise
```

Configuration for unicast streaming video

Procedure steps

1. Specify Module QoS configuration

```
configure terminal
qos module
```

2. Configure a policy map and class for Video

```
policy-map enterprise
class-map video
```

3. Configure rules to classify video packets (assuming port 60000 is used for streaming)

```
match ipv4 protocol 17 dest-port 60000
```

4. Mark video traffic with AF31 DSCP so that other succeeding devices provide the right QoS

```
mark-dscp 26
```

5. Assign a WRR queue with higher weight to the video traffic to give better than best effort service (by default the packet processor is configured to come up with 2 wrr groups (5,6 & 7, 8) in decreasing order of weight)

```
assign-queue 5
```

6. Police video traffic to 400Kbps and assign drop precedence to packets based on conformance levels (by default the packet processor is programmed to come up with stringent WRED thresholds for higher drop precedence)

```
police
srtcm 400 cbs 3000 ebs 6000
remark-cos
exit
pop
```

7. Shape traffic to streamline bursts

```
interface ethernet 6/1
qos module
queue 6
shape 400 3000
```

```
exit
pop
```

8. Apply the policy map to applicable interfaces

```
interface ethernet 6/1
qos module
service-policy input enterprise
```

Configuration for FTP traffic

Procedure steps

1. Specify Module QoS configuration

```
configure terminal
qos module
```

2. Configure a policy map and class for FTP

```
policy-map enterprise
class-map ftp
```

3. Configure rules to classify FTP packets

```
match ipv4 protocol 6 dest-port 21
match ipv4 protocol 6 src-port 21
```

4. Mark FTP traffic with AF31 DSCP so that other succeeding devices provide the right QoS

```
mark-dscp 0
```

5. Assign a WRR queue with lowest weight to the FTP traffic to give best effort service

```
assign-queue 8
pop
```

6. Apply the policy map to applicable interfaces

```
interface ethernet 6/15
qos module
service-policy input enterprise
```

SLA configuration example

With SLA, if the variable that is monitored exceeds the threshold values, then an action can be executed. The action can be a warning message in the console, SNMP trap message to the server, or syslog message the syslog server.

For example, to display an alert message when the response time of a packet is less than 5ms or greater than 10ms, the following configurations can be used.

Procedure steps

To configure an alert message when the response time of a packet is less than 5ms or greater than 10ms, enter:

```
configure terminal
sla profile 1
udp-echo 10.1.1.2 port 7
action response-time console-logging
threshold-type immediate
threshold-value 5 10
```

In the **action** command **response-time** specifies the SLA parameter to be monitored and **console-logging** specifies the action to be taken.

Another scenario is monitoring one-way jitter between source and destination, with the transmission of a trap message when the average value for 10 packets exceeds 30ms. The configuration for this scenario is as follows.

Procedure steps

To configure one-way jitter, enter:

```
configure terminal
configure# sla profile 1
udp-v6-jitter jitter 2001::3
action jitter-average-src-dest trap
threshold-type average
threshold-value 30 10
```

Chapter 11: Auto QoS policies: default configuration

In either Chassis QoS or Ethernet module QoS, you can modify the default Auto Qos policy maps. However, if you do change the default configuration, to revert to the original configuration for the Auto QoS policy map, you have only one option, which is to reconfigure the modified policy map.

The following sections provide the default Auto QoS profile configurations which you can use as the model to reconfigure the Auto QoS policy maps to their original values.

Chassis Auto QoS policy

```
qos
chassis
policy-map AutoQoSPolicyOut
class-map critical root
cbq cr-percent 10.00000 pr-percent 100.0000 priority 1
match dscp cs7
enable-red ds-red
exit class-map
class-map control root
cbq cr-percent 10.00000 pr-percent 100.0000 priority 2
match dscp cs6
enable-red ds-red
exit class-map
class-map voice root
cbq cr-percent 35.00000 pr-percent 100.0000 priority 3
match dscp cs5
match dscp ef
enable-red ds-red
exit class-map
class-map video root
cbq cr-percent 10.00000 pr-percent 100.0000 priority 4
match dscp cs4
match dscp af41
match dscp af42
match dscp af43
enable-red ds-red
exit class-map
class-map streaming root
cbq cr-percent 10.00000 pr-percent 100.0000 priority 4
match dscp cs3
match dscp af31
match dscp af32
match dscp af33
enable-red ds-red
exit class-map
class-map transaction root
cbq cr-percent 5.000000 pr-percent 100.0000 priority 5
```

```
match dscp cs2
match dscp af21
match dscp af22
match dscp af23
enable-red ds-red
exit class-map class-map
oam root cbq cr-percent 10.00000 pr-percent 100.0000 priority 5
match dscp cs1
match dscp af11
match dscp af12
match dscp af13
enable-red ds-red
exit class-map
class-map best-effort root
cbq cr-percent 9.000000 pr-percent 100.0000
match dscp cs0
match dscp default
enable-red ds-red
exit class-map
exit policy-map
policy-map AutoQoSPolicyOut4Q
class-map network root
cbq cr-percent 20.00000 pr-percent 100.0000 priority 1
match dscp cs6
match dscp cs7
enable-red ds-red
exit class-map
class-map voice root
cbq cr-percent 35.00000 pr-percent 100.0000 priority 2
match dscp cs5
match dscp ef
enable-red ds-red
exit class-map
class-map oam root cbq cr-percent 10.00000 pr-percent 100.0000 priority 4
match dscp cs4
match dscp af41
match dscp af42
match dscp af43
enable-red ds-red
exit class-map
class-map best-effort root
cbq cr-percent 34.00000 pr-percent 100.0000 priority 4
match dscp cs0
match dscp cs1
match dscp af11
match dscp af12
match dscp af13
match dscp cs2
match dscp af21
match dscp af22
match dscp af23
match dscp cs3
match dscp af31
match dscp af32
match dscp af33
match dscp default
enable-red ds-red
exit class-map
exit policy-map
policy-map AutoQoSPolicyIn
class-map critical root
police
srtcm cir-percentage 1.000000
color-aware
exit police
```

```
match dscp cs7
exit class-map
class-map control root
police
srtcm cir-percentage 5.000000
color-aware
exit police
match dscp cs6
exit class-map
class-map voice root
police
srtcm cir-percentage 50.00000
color-aware
exit police
match dscp cs5
match dscp ef
exit class-map
class-map video root
police trtcm cir-percentage 10.00000 pir-percentage 15.00000
color-aware
exit police
match dscp cs4
match dscp af41
match dscp af42
match dscp af43
exit class-map
class-map streaming root
police
srtcm cir-percentage 10.00000
color-aware
exit police
match dscp cs3
match dscp af31
match dscp af32
match dscp af33
exit class-map
class-map transaction root
police
srtcm cir-percentage 5.000000
color-aware
exit police
match dscp cs2
match dscp af21
match dscp af22
match dscp af23
exit class-map
class-map oam root
police
srtcm cir-percentage 15.00000
color-aware
exit police
match dscp cs1
match dscp af11
match dscp af12
match dscp af13
exit class-map
class-map best-effort root
match dscp cs0
match dscp default
exit class-map
exit policy-map
```

Ethernet module Auto QoS policy

```
qos
module
policy-map AutoQoSPolicy
class-map critical
police
exit police
match ipv4 dscp cs7
match ipv6 dscp cs7
assign-queue 1
accounting enable
exit class-map
class-map control
police
exit police
match ipv4 dscp cs6
match ipv6 dscp cs6
assign-queue 2
accounting enable
exit class-map
class-map voice
police
exit police
match ipv4 dscp cs5
match ipv4 dscp ef
match ipv6 dscp cs5
match ipv6 dscp ef
assign-queue 3
accounting enable
exit class-map
class-map video
police
exit police
match ipv4 dscp cs4
match ipv4 dscp af41
match ipv4 dscp af42
match ipv4 dscp af43
match ipv6 dscp cs4
match ipv6 dscp af41
match ipv6 dscp af42
match ipv6 dscp af43
assign-queue 4
accounting enable
exit class-map
class-map streaming
police
exit police
match ipv4 dscp cs3
match ipv4 dscp af31
match ipv4 dscp af32
match ipv4 dscp af33
match ipv6 dscp cs3
match ipv6 dscp af31
match ipv6 dscp af32
match ipv6 dscp af33
assign-queue 5
accounting enable
exit class-map
```

```
class-map transaction
police
exit police
match ipv4 dscp cs2
match ipv4 dscp af21
match ipv4 dscp af22
match ipv4 dscp af23
match ipv6 dscp cs2
match ipv6 dscp af21
match ipv6 dscp af22
match ipv6 dscp af23
assign-queue 6
accounting enable
exit class-map
class-map oam
police
exit police
match ipv4 dscp cs1
match ipv4 dscp af11
match ipv4 dscp af12
match ipv4 dscp af13
match ipv6 dscp cs1
match ipv6 dscp af11
match ipv6 dscp af12
match ipv6 dscp af13
assign-queue 7
accounting enable
exit class-map
class-map best-effort
police
exit police
match ipv4 dscp cs0
match ipv6 dscp cs0
assign-queue 8
accounting enable
exit class-map
exit policy-map
exit module
```

Comments? infodev@avaya.com