

Quality of Service—Modular Quality of Service Command Line Interface Single Source Migration Strategy

Last updated: February 2008

Summary

Cisco IOS® Software Release 12.0(26)S supports the Modular Quality-of-Service (QoS) Command-Line-Interface (MQC). It will be added to Release 12.2 (RIs7)S. Releases 12.2 and 12.2T have historically supported MQC differently for distributed hardware (for example, Cisco® 7500 Series Routers and FlexWAN) than for nondistributed hardware (for example, Cisco 7200 Series Routers and other lower-end hardware). Cisco IOS Software will provide a consistent and common MQC behavior on both distributed and nondistributed hardware. Certain MQC behaviors that appear in Cisco 7200 Series Router and other lower-end hardware feature sets will be eliminated or replaced with the behavior from the Cisco 7500 Series Router implementation.

The resulting common behavior has been applied to Releases 12.0(26)S and later, future 12.2(RIs7)S and T train releases. This document describes the new behavioral changes and CLI modifications to the QoS features on the nondistributed hardware (Cisco 7200 Series Router and other lower-end hardware) that result from this effort.

Document Purpose

The purpose of this document is to help service providers that might be affected by the behavioral changes to migrate from Release 12.2T to Release 12.2S. Enterprise customers will be affected when these changes are ported to future T train release. Customers migrating to releases 12.0 (26)S and later will not be affected, because the QoS functionality support was introduced in this release.

This document is not intended to discuss all of the new MQC feature capabilities that have been introduced on the Cisco 7200 Series Router and other lower-end hardware. For a complete list of new feature capabilities, please refer to the appropriate Cisco IOS Software release documentation on Cisco.com.

Benefits of Common Behavior

- Feature parity for Releases 12.2S and 12.0(26)S with future T train releases
- Instant CLI and semantics consistency of QoS features across all hardware that supports Cisco IOS Software
- Common functionality for both distributed and nondistributed implementations, providing consistency of QoS feature behavior across all software-forwarding hardware
- Behavioral consistency across hardware, resulting in accelerated delivery of feature enhancements and new QoS features in different Cisco IOS Software releases
- All hierarchical scheduling capabilities (for example, flow-based fair queuing within a class, hierarchical shaping, and queuing) of the Cisco 7500 Series Router will also be available on the Cisco 7200 Series Router and other lower-end hardware

Effect of QoS Features/CLI Changes on Customers

- Certain QoS feature behaviors will be affected on the Cisco 7200 Series Router and other lower-end hardware.
- To help ensure easy customer migration, certain default values and default feature behaviors will remain unchanged on the Cisco 7200 Series Router and lower-end hardware. Some default parameters might be different across hardware because of various hardware dependencies.
- Certain output commands will be obsolete and no longer supported on the Cisco 7200 Series Router and lower-end hardware. The CLI output for the show policy-map interface command will have some minor changes.
- Certain new QoS feature capabilities will be introduced on both the Cisco 7200 and 7500 Series Routers.

The rest of this document describes the behavioral and CLI changes in more detail, provides any workarounds necessary to achieve the old behavior, and discusses possible side effects to the new behavior.

QoS Feature Behavioral Changes

Fair Queuing

Old Behavior

The fair queuing feature uses a weighted fair queuing (WFQ) function with weights based on IP precedence value and the volume of the flow to make bandwidth guarantees to different flows within a class.

New Behavior

The fair queuing feature is flow based and uses an unweighted queuing function (that is all flows are given equal share of the available bandwidth irrespective of their IP precedence values).

Possible Side Effects of New Behavior

Cisco Systems® recommends that users have separate user-defined classes for traffic that needs to meet certain application requirements/criteria (for example, throughput, latency, and so on). Performance problems under the new behavior might occur if the QoS deployment relies primarily on WFQ in the class default, rather than allocating separate resources for traffic that needs additional help.

Workaround

As of July 2005, there is no compelling technical reason to mandate WFQ in the default class, because MQC has been successfully deployed on the Cisco 7500 Series Router.

Allocation of Bandwidth to Class Default

Old Behavior

The default class can use up to 25 percent of total available bandwidth; however, the entire 25 percent is not guaranteed. Rather, it is proportionately shared between different flows in the default class and excess traffic from other bandwidth classes. Thus, the amount of bandwidth that the default class will receive depends on a number of factors, including the total number of flows currently in the router, the bandwidth guarantees (or weights) made to the other user-defined classes, and the number of hash queues in the router. To make minimum bandwidth guarantees to the default class, the bandwidth command needs to be explicitly configured under the class in the policy.

New Behavior

The class default has a default minimum guarantee that equals the difference between the total available bandwidth (for example, link rate, shaped rate) and the amount of bandwidth guaranteed to the other classes. For example, if 90 percent of bandwidth is allocated to other classes, then the class default is guaranteed the remaining 10 percent. If there is no traffic in the class default, then the other classes share that 10 percent proportionally. Alternatively, the user can explicitly configure the amount of bandwidth that should be available to default class using the bandwidth <x> command. This will lower the guarantee that is given to the class default and allow 10 minus “x” to always be available for the other classes.

Workaround

As of July 2005, there is no workaround available.

Class-Based Shaping Support on ATM Permanent Virtual Circuits

Old Behavior

Class-based traffic shaping is not supported in ATM interfaces, because traffic shaping is native to the ATM framework. To achieve different levels of QoS, users must configure the permanent virtual circuits (PVCs) as either real-time variable bit rate (VBR-rt), non-real-time variable bit rate (VBR-nrt), constant bit rate (CBR), available bit rate (ABR), or unspecified bit rate (UBR).

New Behavior

Class-based traffic shaping is supported on all ATM interfaces. This will allow an extra level of specificity by allowing users to shape one or more classes under the ATM traffic contract.

Hierarchical Queuing Support

Old Behavior

Support exists for hierarchical queuing/policies only if shaping is configured at the parent policy.

New Behavior

Hierarchical queuing also supports the **bandwidth** command configured at the parent policy. The following sample configuration illustrates the new hierarchical queuing with the **bandwidth** command configured at the parent level:

```
policy child-foo
  class child1-c1
    bandwidth 40
```

```
class child2-c1
    bandwidth 60

policy foo
    class c1
        bandwidth 100
    service-policy child-foo
```

Changes in Default Values and Behavior

Default Queuing Strategy for Class Default Changes from WFQ to First In, First Out

Old Behavior

The queuing strategy is weighted fair between the different flows in the class default. If any bandwidth guarantee to default class is made using the **bandwidth** statement, queuing within the default class is no longer WFQ between the different flows, but simply first in, first out (FIFO). Class-based shaping also has implicit fair-queuing behavior within the class.

New Behavior

The class-default traffic goes into a FIFO queue, and by default there is no fair queuing involved. Class-based shaping will have implicit FIFO instead of fair-queuing behavior.

Workaround

The user can explicitly configure fair queuing in this class to create a layer of flow queues within the FIFO queue in the class default.

Dual FIFO Behavior for Frame Relay Traffic Shaping

Old Behavior

When FRF.12 fragmentation is configured on a Frame Relay PVC, the interface queue becomes a dual FIFO queue. Low Latency Queuing (LLQ) traffic from all the virtual circuits goes to the high-priority FIFO queue, while traffic from the remaining bandwidth classes goes to the low-priority FIFO queue. The low-priority queue is serviced only when the high-priority queue is empty. This is done to meet latency guarantees for LLQ traffic under the previous scheduler implementation.

New Behavior

There will be no interface-level dual FIFO queuing support for Frame Relay traffic shaping (FRTS). The **show interface** command will also not display dual FIFO as a queuing strategy. The functionality is removed because it is no longer needed to meet performance expectations under the new behavior.

Weighted Random Early Detection Minimum/Maximum Threshold Changes

Old Behavior

The default min/max threshold values for precedence/Differentiated Services Code Point (DSCP)-based Weighted Random Early Detection (WRED) are based on a fixed value table for the Cisco 7200 Series Router and lower-end hardware. Users can change the default threshold settings by using the **random-detect dscp <dscp> <min threshold> <max threshold>** command.

New Behavior

The old behavior will be maintained for DSCP and IP precedence-based WRED on the Cisco 7200 Series Router and lower-end hardware.

The new byte-based and time-based WRED features introduced in Cisco IOS Software Release 12.2S will have their default min/max threshold values derived from the available system buffer size.

Obsolete and New Commands

Support for Fair-Queue Queue-Limit Command

Old Behavior

None.

New Behavior

Users can use the **fair-queue queue-limit < >** command to set per flow queue limit for each dynamic flow queue within a class.

Shape Max-Buffers

Old Behavior

This command specifies the maximum number of buffers allowed in each shaping queue.

New Behavior

Not supported.

Workaround

The **queue-limit** command should be used to specify the depth of the shaping queue. This command is allowed on any class that has **bandwidth** or **shape** enabled.

Max-Reserved-Bandwidth

Old Behavior

The default maximum reserved bandwidth is 75 percent, so the maximum bandwidth that can be guaranteed to any user-defined class is also 75 percent. If 75 percent of the bandwidth is allocated only for the LLQ, then no minimum bandwidth can be guaranteed to the other classes, and they will share the remaining 25 percent bandwidth with the class default traffic.

If more bandwidth needs to be allocated, use the **max-reserved-bandwidth** command to modify the bandwidth amount that can be reserved for user-defined classes.

New Behavior

The **max-reserved-bandwidth** command no longer affects the amount of bandwidth available to a service-policy. 1% must be reserved for the class-default with the rest being available to the users classes. Please also refer to the previous section "Allocation of Bandwidth to Class Default."

CLI Output Changes

"Show Queuing" and "Show Queue" Commands

The above two commands will no longer be supported. Instead, users can use the **show policy-map interface** and **show policy-map** commands to gather QoS-related information and statistics.

“Show Policy-Map Interface” and “Show Policy-Map” Output

The new **show policy-map interface** output has additional fields that display the DSCP value, WRED statistics in bytes, new column for number of transmitted packets by WRED, and a counter displaying packets output/bytes output in each class. The **packets queued/bytes queued** counter will no longer be supported.

Table 1 compares the old and new output formats for both commands.

Table 1. Comparison of Old and New Command Output Formats

Comments	New Output Format	Old Output Format
	Show policy-map parent Policy Map parent Class class-default Average Rate Traffic Shaping cir 300000 (bps) bc 1200 (bits) be 1200 (bits) service-policy child Show policy-map child Policy Map child Class prec2 priority 40 (%) Class prec4 bandwidth 40 (%) packet-based wred, exponential weight 9 class min-threshold max-threshold mark-probability	Show policy-map parent Policy Map parent Class class-default Traffic Shaping Average Rate Traffic Shaping CIR 300000 (bps) Max. Buffers Limit 1000 (Packets) Bc 1200 Be 1200 service-policy child Show policy-map child Policy Map child Class prec2 Bandwidth 40 (%) Max Threshold 64 (packets) Class prec4 Bandwidth 40 (%) exponential weight 9 class min-threshold max-threshold mark-probability
	0 - - 1/10 1 - - 1/10 2 - - 1/10 3 - - 1/10 4 - - 1/10 5 - - 1/10 6 - - 1/10 7 - - 1/10 Class class-default fair-queue packet-based wred, exponential weight 9 class min-threshold max-threshold mark-probability	0 - - 1/10 1 - - 1/10 2 - - 1/10 3 - - 1/10 4 - - 1/10 5 - - 1/10 6 - - 1/10 7 - - 1/10 rsvp - - 1/10 Class class-default Flow based Fair Queueing exponential weight 9 dscp min-threshold max-threshold mark-probability

	<pre> 0 -- 1/10 1 -- 1/10 2 -- 1/10 3 -- 1/10 4 -- 1/10 5 -- 1/10 6 -- 1/10 7 -- 1/10 </pre>	<pre> af11 -- 1/10 af12 -- 1/10 af13 -- 1/10 af21 -- 1/10 af22 -- 1/10 af23 -- 1/10 af31 -- 1/10 af32 -- 1/10 af33 -- 1/10 af41 -- 1/10 af42 -- 1/10 af43 -- 1/10 cs1 25 50 1/10 cs2 -- 1/10 cs3 -- 1/10 cs4 -- 1/10 cs5 -- 1/10 cs6 -- 1/10 cs7 -- 1/10 ef -- 1/10 rsvp -- 1/10 default -- 1/10 </pre>
	<pre> show policy-map int fas0/0 FastEthernet0/0 Service-policy output: parent (1045) Class-map: class-default (match-any) (1058/0) 414491 packets, 24869460 bytes 30 second offered rate 3571000 bps, drop rate 3495000 bps Match: any (1059) 414491 packets, 24869460 bytes 30 second rate 3571000 bps Queueing queue limit 64 packets (queue depth/total drops/no-buffer drops) 40/405784/0 (pkts queued/bytes queued) 8768/526080 shape (average) cir 300000, bc 1200, be 1200 target shape rate 300000 Service-policy : child (1046) queue stats for all priority classes: queue limit 64 packets (queue depth/total drops/no-buffer drops) 0/203740/0 (pkts queued/bytes queued) 4180/250800 Class-map: prec2 (match-any) (1047/2) 207246 packets, 12434760 bytes 30 second offered rate 1787000 bps, drop rate 1754000 bps </pre>	<pre> Show policy-map interface FastEthernet5/0 Service-policy output: parent Class-map: class-default (match-any) 218251 packets, 20951766 bytes 30 second offered rate 3811000 bps, drop rate 3706000 bps Match: any Traffic Shaping Target/Average Byte Sustain Excess Interval Increment Rate Limit bits/int bits/int (ms) (bytes) 300000/300000 300 1200 1200 4 150 Adapt Queue Packets Bytes Packets Bytes Shaping Active Depth Delayed Delayed Active - 0 6153 590358 6143 589508 no Service-policy : child Class-map: prec2 (match-all) 72637 packets, 6973152 bytes 30 second offered rate 1270000 bps, drop rate 1218000 bps </pre>

<p>Match: ip precedence 2 (1048) 207246 packets, 12434760 bytes 30 second rate 1787000 bps Priority: 40% (120 kbps), burst bytes 3000, Class-map: prec4 (match-any) (1050/1) 207247 packets, 12434820 bytes 30 second offered rate 1787000 bps, drop rate 1741000 bps Match: ip precedence 4 (1051) 207247 packets, 12434820 bytes 30 second rate 1787000 bps Queueing queue limit 64 packets (queue depth/total drops/no-buffer drops) 35/238952/0 (pkts queued/bytes queued) 6172/370320 bandwidth 40% (120 kbps) Exp-weight-constant: 9 (1/512) Mean queue depth: 33 packets class Transmitted Random drop Tail drop Minimum Maximum Mark</p>	<p>Match: ip precedence 2 Queueing Output Queue: Conversation 41 Bandwidth 40 (%) Max Threshold 64 (packets) (pkts matched/bytes matched) 72635/6972960 (depth/total drops/no-buffer drops) 0/69643/0 Class-map: prec4 (match-all) 72880 packets, 6996480 bytes 30 second offered rate 1276000 bps, drop rate 1224000 bps Match: ip precedence 4 Queueing Output Queue: Conversation 42 Bandwidth 40 (%) (pkts matched/bytes matched) 72878/6996288 (depth/total drops/no-buffer drops) 0/69909/0 exponential weight: 9 mean queue depth: 0 class Transmitted Random drop Tail drop Minimum Maximum Mark</p>
<p>pkts/bytes pkts/bytes pkts/bytes thresh thresh prob 0 0/0 0/0 0/0 16 32 1/10 1 0/0 0/0 0/0 18 32 1/10 2 0/0 0/0 0/0 20 32 1/10 3 0/0 0/0 0/0 22 32 1/10 4 5219/313140 522/31320 201500/12090000 24 32 1/10 5 0/0 0/0 0/0 26 32 1/10 6 0/0 0/0 0/0 28 32 1/10 7 0/0 0/0 0/0 30 32 1/10 Class-map: class-default (match-any) (1054/0) 3 packets, 180 bytes 30 second offered rate 0 bps, drop rate 0 bps Match: any (1055) 3 packets, 180 bytes 30 second rate 0 bps Queueing queue limit 64 packets (queue depth/total drops/no-buffer drops/flowdrops) 0/0/0/0 (pkts queued/bytes queued) 4/240 Fair-queue: per-flow queue limit 16 Exp-weight-constant: 9 (1/512) Mean queue depth: 0 packets class Transmitted Random drop Tail/Flow drop Minimum Maximum Mark</p>	<p>pkts/bytes pkts/bytes pkts/bytes thresh thresh prob 0 0/0 0/0 0/0 20 40 1/10 1 0/0 0/0 0/0 22 40 1/10 2 0/0 0/0 0/0 24 40 1/10 3 0/0 0/0 0/0 26 40 1/10 4 2971/285216 501/48096 69408/6663168 28 40 1/10 5 0/0 0/0 0/0 30 40 1/10 6 0/0 0/0 0/0 32 40 1/10 7 0/0 0/0 0/0 34 40 1/10 rsvp 0/0 0/0 0/0 36 40 1/10 Class-map: class-default (match-any) 0 packets, 0 bytes 30 second offered rate 0 bps, drop rate 0 bps Match: any Queueing Flow Based Fair Queueing Maximum Number of Hashed Queues 32 November 25, 2003 QoS Queueing code merge Functional Spec.Rev A (total queued/total drops/no-buffer drops) 0/0/0 exponential weight: 9 dscp Transmitted Random drop Tail drop Minimum Maximum Mark</p>

	pkts/bytes	pkts/bytes	pkts/bytes	thresh	thresh	prob	pkts/bytes	pkts/bytes	pkts/bytes	thresh	thresh	prob
	0	0/0	0/0	16	32	1/10	af11	0/0	0/0	32	40	1/10
	1	0/0	0/0	18	32	1/10	af12	0/0	0/0	28	40	1/10
	2	0/0	0/0	20	32	1/10	af13	0/0	0/0	24	40	1/10
	3	0/0	0/0	22	32	1/10	af21	0/0	0/0	32	40	1/10
	4	0/0	0/0	24	32	1/10	af22	0/0	0/0	28	40	1/10
	5	0/0	0/0	26	32	1/10	af23	0/0	0/0	24	40	1/10
	6	0/0	0/0	28	32	1/10	af31	0/0	0/0	32	40	1/10
	7	0/0	0/0	30	32	1/10	af32	0/0	0/0	28	40	1/10
							af33	0/0	0/0	24	40	1/10
							af41	0/0	0/0	32	40	1/10
							af42	0/0	0/0	28	40	1/10
							af43	0/0	0/0	24	40	1/10
							cs1	0/0	0/0	25	50	1/10
							cs2	0/0	0/0	24	40	1/10
							cs3	0/0	0/0	26	40	1/10
							cs4	0/0	0/0	28	40	1/10
							cs5	0/0	0/0	30	40	1/10
							cs6	3/222	0/0	32	40	1/10
							cs7	0/0	0/0	34	40	1/10
							ef	0/0	0/0	36	40	1/10
							rsvp	0/0	0/0	36	40	1/10
							default	0/0	0/0	20	40	1/10

References

- [Cisco QoS Features in Cisco IOS Software Release 12.0S](#)
- [Cisco New Features in Cisco IOS Software Release 12.2S](#)
- [Cisco New Features in Cisco IOS Software Release 12.4T](#)



Americas Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
 www.cisco.com
 Tel: 408 526-4000
 800 553-NETS (6387)
 Fax: 408 527-0883

Asia Pacific Headquarters
 Cisco Systems (USA) Pte. Ltd.
 168 Robinson Road
 #28-01 Capital Tower
 Singapore 068912
 www.cisco.com
 Tel: +65 6317 7777
 Fax: +65 6317 7799

Europe Headquarters
 Cisco Systems International BV
 Haarlerbergpark
 Haarlerbergweg 13-19
 1101 CH Amsterdam
 The Netherlands
 www-europe.cisco.com
 Tel: +31 0 800 020 0791
 Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, IQ Expertise, the IQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0711R)