



MQC Hierarchical Queuing with 3 Level Scheduler

First Published: November, 2006

The MQC Hierarchical Queuing with 3 Level Scheduler feature provides a flexible packet scheduling and queuing system in which you can specify how excess bandwidth is to be allocated among the subscriber (logical) queues.

History for the MQC Hierarchical Queuing with 3 Level Scheduler Feature

Release	Modification
12.2(31)SB2	This feature was introduced and implemented on the Cisco 10000 series router for the PRE3.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at <http://www.cisco.com/go/fn>. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.

Contents

- [Prerequisites for the Three-Level Scheduler, page 2](#)
- [Restrictions for the Three-Level Scheduler, page 2](#)
- [Information About the MQC Hierarchical Queuing with 3 Level Scheduler, page 2](#)
- [How to Configure Bandwidth-Remaining Ratios, page 6](#)
- [Configuration Examples for the Three-Level Scheduler, page 6](#)
- [Additional References, page 10](#)
- [Command Reference, page 11](#)



Corporate Headquarters:

Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2006 Cisco Systems, Inc. All rights reserved.

Prerequisites for the Three-Level Scheduler

Traffic classes must be configured on the router using the **class-map** command.

Restrictions for the Three-Level Scheduler

- The priority queue in a child policy must be policed to 90 percent of the parent's shaped bandwidth.
- The three-level scheduler does not support bandwidth propagation. Therefore, you cannot configure a bandwidth guarantee for any queue other than a priority queue.
- To allow oversubscription provisioning, the admission control check is not performed.
- The three-level scheduler does not allocate an implicit bandwidth guarantee for the parent class-default class. Instead, the scheduler uses the ratio of the classes to allocate bandwidth.
- When hierarchical policies are enabled on multiple VLANs and each VLAN hierarchical policy has priority services configured in a child policy, the three-level scheduler first services the priority traffic from all VLANs and then proportionally shares the remaining bandwidth of the interface among all of the VLANs.



Note

The two-level scheduler allocates an equal share of interface bandwidth to each VLAN. After the two-level scheduler serves priority services, best-effort traffic from a VLAN uses the remaining bandwidth. If priority traffic is not configured, instead of proportionally allocating the remaining bandwidth available to each VLAN, the two-level scheduler allocates the entire interface bandwidth to the VLAN's best-effort traffic.

- The sum of all priority traffic running on a given port must be less than or equal to 90 percent of the port bandwidth.

Information About the MQC Hierarchical Queuing with 3 Level Scheduler

The MQC Hierarchical Queuing with 3 Level Scheduler feature provides a flexible packet scheduling and queuing system in which you can specify how excess bandwidth is to be allocated among the subscriber queues and logical interfaces. Rather than allocating an implicit minimum bandwidth guarantee to each queue, the three-level scheduler uses the bandwidth-remaining ratio parameter to allocate unused bandwidth to each logical queue. The three-level scheduler services queues based on the following user-configurable parameters:

- Maximum rate—The specified shape rate of the parent queue.
- Bandwidth-remaining ratio—The value used to determine the portion of unused, non-guaranteed bandwidth allocated to a logical queue relative to other queues competing for the unused bandwidth.



Note

At the class level, the router converts the values specified in the **bandwidth bps** and **bandwidth remaining percent** commands to a bandwidth-remaining ratio value. The router does not allow you to configure the **bandwidth bps** and **bandwidth remaining percent** commands on the physical and logical layers.

The three-level scheduler on the PRE3 supports priority propagation by propagating the priority guarantees you configure for subscriber services down to the logical interface level. Therefore, the priority traffic is serviced first at the logical and class level. After servicing the priority traffic bandwidth, the three-level scheduler allocates unused bandwidth to the logical queues based on the configured bandwidth-remaining ratio. In this default case, the three-level scheduler allocates an equal share of the unused bandwidth to each logical queue.

The three-level scheduler supports shaping and scheduling only on the egress interface. The **bandwidth** command must be configured as a percentage of the available bandwidth or as an absolute bandwidth. You cannot concurrently configure the **bandwidth** and **bandwidth remaining** commands on the same class queue or the same policy map.

For more information about the bandwidth-remaining ratio, see the *Distribution of Remaining Bandwidth Using Ratio feature module*.

Modular QoS Command-Line Interface

The Modular Quality of Service Command-Line Interface (MQC) is designed to simplify the configuration of Quality of Service (QoS) on Cisco routers and switches by defining a common command syntax and resulting set of QoS behaviors across platforms. This model replaces the previous model of defining unique syntaxes for each QoS feature and for each platform.

The MQC contains the following three steps:

- Define a traffic class using the **class-map** command.
- Create a traffic policy by associating the traffic class with one or more QoS features using the **policy-map** command.
- Attach the traffic policy to the interface, subinterface, or virtual circuit (VC) using the **service-policy** command.

For more information about MQC, see the [Modular Quality of Service Command-Line Interface](http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5014/products_feature_guide_book09186a0080088141.html) document at the following URL:

http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5014/products_feature_guide_book09186a0080088141.html

Scheduling Hierarchy

As shown in [Figure 1 on page 4](#), the three-level scheduler uses the following scheduling hierarchy to allocate bandwidth for subscriber traffic:

- Class layer—The three-level scheduler uses virtual-time calendars to schedule class queues and logical interfaces.
- Logical layer (VLAN or ATM VC)—Virtual-time calendars perform weighted round robin based on the weight of the logical interface and the number of bytes dequeued.
- Physical layer (interface or ATM virtual path)—Token buckets ensure that the maximum rate for the class and the logical interface are not exceeded.

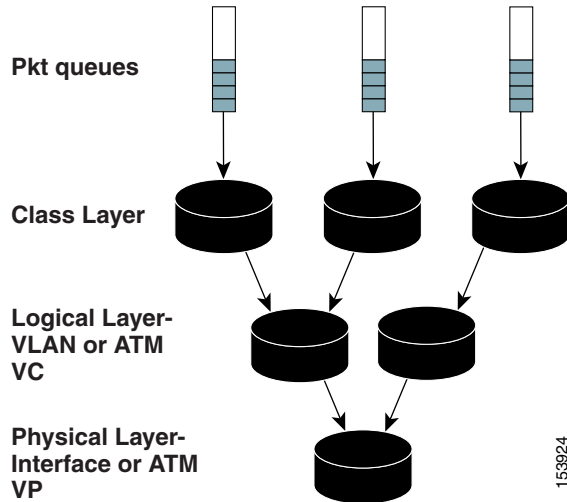
Figure 1 **Scheduling Hierarchy**

Table 1 provides an example of how the scheduling hierarchy can apply to Ethernet and ATM topologies. For Ethernet, you cannot oversubscribe the Queue-in-Queue (qinq) into the interface. For ATM, you cannot oversubscribe the virtual path (VP) into the interface.

Table 1 **Applying the Scheduling Hierarchy to Ethernet and ATM**

Scheduling Hierarchy	Ethernet	ATM
Class layer (virtual time)	MQC-defined queues	MQC-defined queues
Logical layer (virtual time)	VLAN (inner tag)	Virtual channel (VC)
	Session	
Physical (real time)	Queue-in-Queue (outer tag)	Virtual path (VP)
	VLAN (inner tag), if session is the logical layer identifier	

By using VP and VC scheduling with existing Cisco 10000 ATM line cards, the scheduler supports priority propagation: cell-based VP shaping in the segmentation and reassembly (SAR) mechanism with frame-based VC scheduling in the performance routing engine 3 (PRE3).

Priority Service and Latency

The three-level scheduler supports multiple levels of priority service that you can use for such purposes as control traffic, delay-sensitive traffic (for example, voice), minimum guarantees, and excess bandwidth allocation. Each level of priority supports multiple queues, which allows for multiple types of delay-sensitive traffic (for example, voice and video).

The three-level scheduler can service the same queue from multiple levels of priority service. For example, the three-level scheduler uses priority level 1 for voice, priority level 2 for video, and the excess bandwidth for data.

For a priority class with policing configured, the three-level scheduler always polices the priority traffic to the rate specified in the **police** command (1000 kbps as shown in the following example configuration), regardless of whether or not the underlying interface is congested.

```
Router(config-pmap-c)# priority
Router(config-pmap-c)# police 1000
```

**Note**

The three-level scheduler does not support the **priority kbps** command.

Latency Requirements

Delay-sensitive traffic incurs a maximum of 10 milliseconds (ms) of latency on edge router interfaces and a maximum of 1 ms of latency on core router interfaces. For interface speeds at T1/E1 and below, the three-level scheduler services 2 maximum transmission units (MTUs) of nonpriority traffic before servicing a priority packet. Requirements for high-speed interfaces are not as strict as 2 MTUs, but are always bound by 10 ms on edge interfaces and 1 ms on core interfaces.

The three-level scheduler also supports the minimal latency requirement (2 MTUs of nonpriority traffic in front of priority traffic) at the physical link rate. However, in some cases, it is impossible for the three-level scheduler to service all competing packets with a latency of 2 MTUs. For example, if many priority packets compete at the same time for bandwidth, the last one serviced may incur latency that is greater than 2 MTUs.

[Table 2](#) lists the maximum latency requirements for various interface speeds.

Table 2 **Maximum Latency Requirements**

Interface Speed	Maximum Latency
Greater than 2 Mbps	2 MTU + 6 ms
2 Mbps to 1 Gbps	2 MTU
1 Gbps or greater	1 ms

Priority Propagation with Imposed Burstiness

A single physical interface can have large numbers of logical interfaces and each of these logical interfaces can have both priority and nonpriority traffic competing for the physical link. To minimize latency, the priority traffic of one logical interface has priority over the nonpriority traffic of other logical interfaces, thereby imposing burstiness on the minimum rate traffic of other logical interfaces. The latency that the priority traffic incurs results from the rate constraining the delivered rate of the priority traffic. In many cases, this constraining rate is not the rate of the priority class's parent policy.

For example, suppose a 10 Gigabit Ethernet (GE) interface has 100 VLANs that are shaped to various rates. Each VLAN has a priority class and additional classes configured. Through priority propagation, the scheduler delivers latency to the priority traffic based on the 10 GE rate and not the VLAN rate.

**Note**

The VLAN rate is at most 1 to 2 MTUs of nonpriority traffic in front of priority traffic, which would bound the latency incurred by priority traffic (due to non-priority traffic) at 1 to 2 MTUs served at the 10 GE rate.

The priority traffic of one logical interface cannot only impose burstiness on other traffic, but also starve other traffic. The only way to prevent the starvation of other traffic is by configuring a policer on the priority queue by limiting the percent of priority traffic to less than 90 percent of the parent bandwidth and the port bandwidth.

Configuration Granularity

Table 3 describes the configuration granularity for the three-level scheduler.

Table 3 Three-Level Scheduler Configuration Granularity

Interface Bandwidth	Granularity
Less than or equal to 2 Mbps	.4%
Greater than 2 Mbps and less than 1 Gbps	.2%
Greater than or equal to 1 Gbps	.1%

How to Configure Bandwidth-Remaining Ratios

To configure bandwidth-remaining ratios on subinterface-level and class-level queues, see the *Distribution of Remaining Bandwidth Using Ratio, Release 12.2(31)SB2 feature module*.

Configuration Examples for the Three-Level Scheduler

This section provides the following configuration examples:

- [Bandwidth Allocation—Policy Attached to an Interface: Example, page 6](#)
- [Bandwidth Allocation—Parent Policy Attached to Two Subinterfaces: Example, page 7](#)
- [Tuning the Bandwidth-Remaining Ratio: Example, page 8](#)

Bandwidth Allocation—Policy Attached to an Interface: Example

The following example configuration consists of one policy map named Child with the following traffic classes defined: prec0, prec2, and class-default. The policy is attached to the ATM interface 1/0/0, which has a configured rate of 1000 kbps.

```
policy-map Child
  class prec0
    bandwidth 300
  class prec2
    bandwidth 100
  class class-default
    bandwidth 50
!
interface atm 1/0/0
  bandwidth 1000
  service-policy output Child
```

Assuming that the traffic flow through each class is enough to require maximum possible bandwidth, the three-level scheduler allocates bandwidth as described in [Table 4](#).

Table 4 *Queuing Presentation—Policy Attached to an Interface*

Traffic Class	Bandwidth Ratio	Total Bandwidth Allocated
prec0	6	666 kbps
prec2	2	222 kbps
class-default	1	111 kbps

Bandwidth Allocation—Parent Policy Attached to Two Subinterfaces: Example

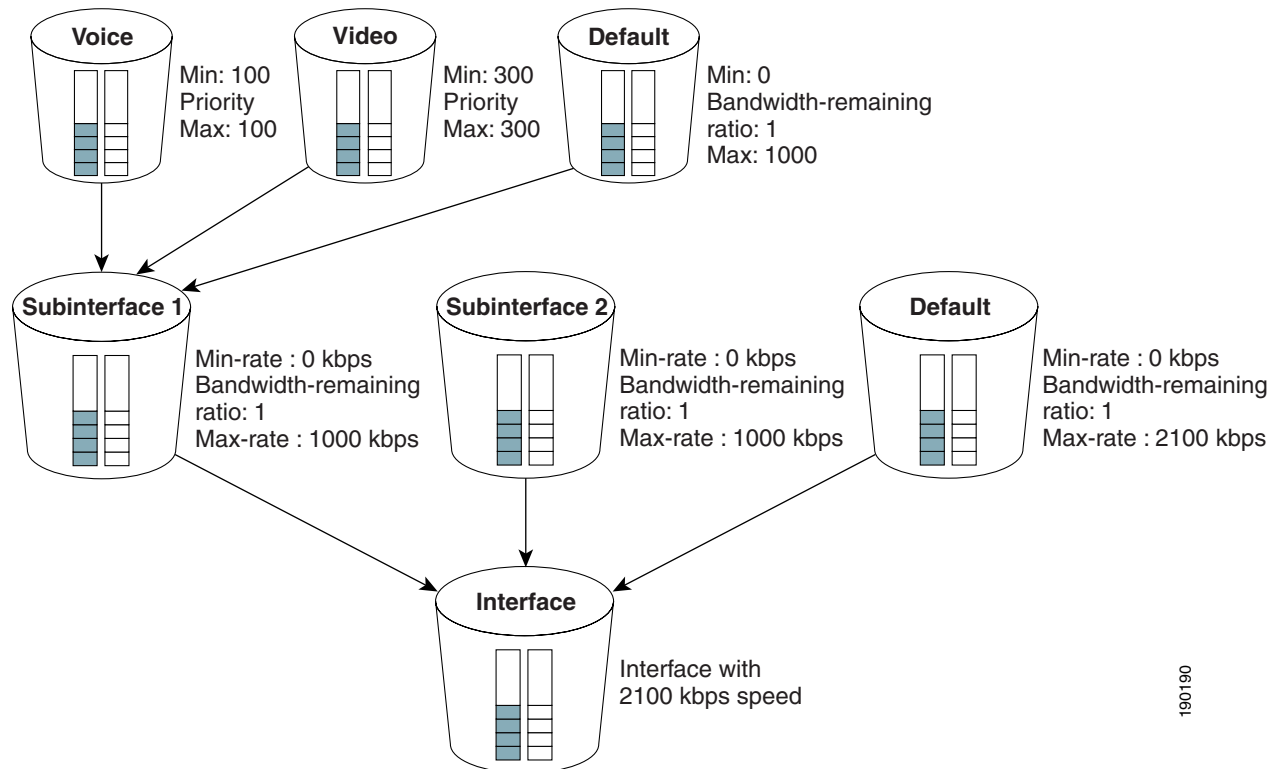
The following example configuration contains a hierarchical policy consisting of two policy maps: Child and Parent. The Child policy has two traffic classes (voice and video) with each configured as a priority class with policing enabled. The Parent policy has its class-default class shaped to 1000 kbps. The Parent policy is attached to the ATM subinterface 1/0/1.1 and to subinterface 1/0/1.2. ATM interface 1/0/1 has a configured rate of 2100 kbps.

```

policy-map Child
  class voice
    priority level 1
    police 100
  !
  class video
    priority level 2
    police 300
  !
policy-map Parent
  class class-default
    shape average 1000
    service-policy Child
  !
interface atm 1/0/1
  atm pvp 1 1400
  !
interface atm 1/0/1.1
  bandwidth remaining ratio 1
  service-policy output Parent
  !
interface atm 1/0/1.2
  bandwidth remaining ratio 1
  service-policy output Parent
  !

```

[Figure 2](#) shows an example of the queuing presentation based on the above configuration. The service rates for all Child classes under each subinterface might differ from the rates shown in [Figure 2](#), depending on the presence or absence of priority propagation and how the class's bandwidth usage is accounted against the Parent queue.

Figure 2 **Queuing Presentation—Parent Enabled on Two Subinterfaces**

Each subinterface receives an equal share of bandwidth. Based on the bandwidth-remaining ratio of 1, each subinterface-level queue receives a rate of 700 kbps (subinterfaces 1 and 2 queues, and default queue at subinterface-level).

- For subinterface 1, assume that only the voice traffic is active. From the 700-kbps bandwidth allocated to subinterface 1, the voice traffic receives a bandwidth rate of 100 kbps and the default traffic receives a rate of 600 kbps.
- For subinterface 2, assume that only the video traffic is active. From the 700-kbps bandwidth allocated to subinterface 2, the video traffic receives a bandwidth rate of 300 kbps and the default traffic receives a rate of 400 kbps.

Tuning the Bandwidth-Remaining Ratio: Example

The following example configuration shows how to tune the bandwidth-remaining ratio using the **bandwidth remaining ratio** command. In the example, the class-default class of Parent1 has a bandwidth-remaining ratio of 9 and the class-default class of Parent2 has a bandwidth-remaining ratio of 7.

```
policy-map Child
  class prec0
    priority level 1
    police 100
  !
  class prec2
    priority level 2
    police 300
  !
```

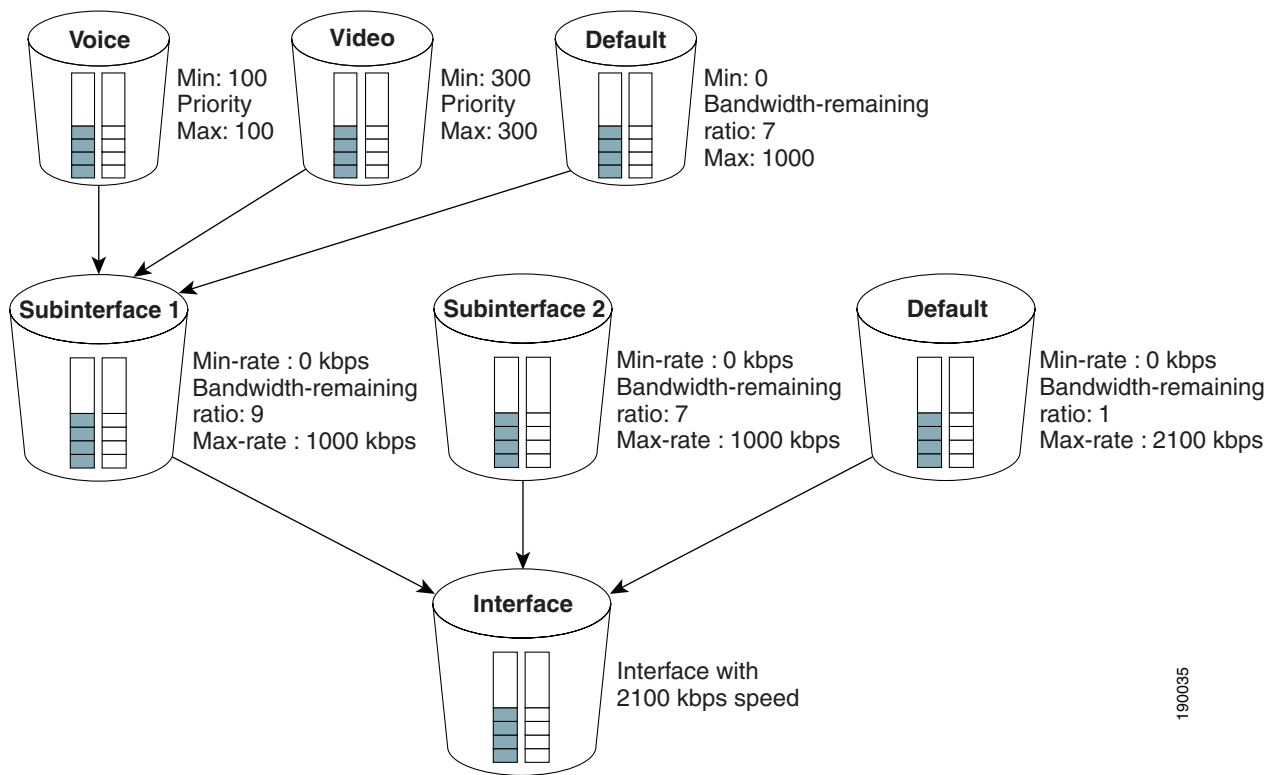
```

policy-map Parent1
  class class-default
    shape average 10000
    bandwidth remaining ratio 9
!
policy-map Parent2
  class class-default
    shape average 1000
    bandwidth remaining ratio 7

```

Figure 3 shows an example of the queuing presentation based on the above configuration and assuming that the Parent1 policy is enabled on subinterface 1 and the Parent2 policy is enabled on subinterface 2, and that the interface speed is 2100 kbps.

Figure 3 Queuing Presentation—Tuning the Bandwidth-Remaining Ratio



Based on the preceding configuration, the three-level scheduler distributes bandwidth in the following way (assuming that the voice traffic is active on subinterface 1 only and the video traffic is active on subinterface 2 only):

- A total of 400 kbps of bandwidth is used from the interface: 100 kbps-bandwidth guarantee for voice traffic on subinterface 1 and 300-kbps bandwidth guarantee for video traffic on subinterface 2.
- The remaining 1700-kbps bandwidth is distributed across the subinterface-level queues based on their bandwidth-remaining ratios:
 - Subinterface 1 with bandwidth-remaining ratio 9 receives 956 kbps
 - Subinterface 2 with bandwidth-remaining ratio 7 receives 743 kbps

Additional References

The following sections provide references related to the MQC Hierarchical Queuing with 3 Level Scheduler feature.

Related Documents

Related Topic	Document Title
Bandwidth	Cisco 10000 Series Router Quality of Service Configuration Guide Distributing Bandwidth Between Queues
Bandwidth-remaining ratio	Distribution of Remaining Bandwidth Using Ratio feature module
Hierarchical policies	Cisco 10000 Series Router Quality of Service Configuration Guide Defining QoS for Multiple Policy Levels
Policy maps	Cisco 10000 Series Router Quality of Service Configuration Guide Configuring QoS Policy Actions and Rules
Shaping traffic	Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2 Part 4: Policing and Shaping > Configuring Class-Based Shaping Part 4: Policing and Shaping > Policing and Shaping Overview > Traffic Shaping > Class-Based Shaping
Traffic policing and shaping	Comparing Traffic Policing and Traffic Shaping for Bandwidth Limiting

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
The Cisco Technical Support & Documentation website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

Command Reference

This section documents new and modified commands only.

- [bandwidth remaining ratio](#)
- [show policy-map](#)
- [show policy-map interface](#)

bandwidth remaining ratio

To specify a bandwidth-remaining ratio for class-level or subinterface-level queues to be used during congestion to determine the amount of excess bandwidth (unused by priority traffic) to allocate to non-priority queues, use the **bandwidth remaining ratio** command in policy-map class configuration mode. To remove the bandwidth-remaining ratio, use the **no** form of this command.

bandwidth remaining ratio *ratio*

no bandwidth remaining ratio *ratio*

Syntax Description

<i>ratio</i>	Specifies the relative weight of this subinterface or queue with respect to other subinterfaces or queues. Valid values are from 1 to 1000. The default value is platform dependent.
--------------	--

Command Default

Cisco 10000 Series Router

When using default bandwidth-remaining ratios at the subinterface level, the Cisco 10000 series router distinguishes between interface types. At the subinterface level, the default bandwidth-remaining ratio is 1 for VLAN subinterfaces and Frame Relay DLCIs. For ATM subinterfaces, the router computes the default bandwidth-remaining ratio based on the subinterface speed.

When using default bandwidth-remaining ratios at the class level, the Cisco 10000 series router makes no distinction between interface types. At the class level, the default bandwidth-remaining ratio is 1.

Command Modes

Policy-map class

Command History

Release	Modification
12.2(31)SB2	This command was introduced and implemented on the Cisco 10000 series router for the PRE3.

Usage Guidelines

Cisco 10000 Series Router

The scheduler uses the ratio specified in the **bandwidth remaining ratio** command to determine the amount of excess bandwidth (unused by priority traffic) to allocate to a class-level queue or a subinterface-level queue during periods of congestion. The scheduler allocates the unused bandwidth relative to other queues or subinterfaces.

The **bandwidth remaining ratio** command cannot coexist with another **bandwidth** command in different traffic classes of the same policy map. For example, the following configuration is not valid and causes an error message to display:

```
policy-map Precl
  class precedence_0
    bandwidth remaining ratio 10
  class precedence_2
    bandwidth 1000
```

For the PRE2, the **bandwidth remaining ratio** command can coexist with another **bandwidth** command in the same class of a policy map. On the PRE3, the **bandwidth remaining ratio** command cannot coexist with another **bandwidth** command in the same class. For example, the following configuration is not valid on the PRE3 and causes an error message to display:

```
policy-map Prec1
  class precedence_0
    bandwidth 1000
    bandwidth remaining ratio 10
```

In a hierarchical policy map in which the parent policy has only the class-default class defined with a child queuing policy applied, the router accepts only the **bandwidth remaining ratio** form of the **bandwidth** command in the class-default class.

The **bandwidth remaining ratio** command cannot coexist with the **priority** command in the same class. For example, the following configuration is not valid and causes an error message to display:

```
policy-map Prec1
  class precedence_1
    priority
    police percent 30
    bandwidth remaining ratio 10
```

All of the queues for which the **bandwidth remaining ratio** command is not specified receive the platform-specified minimum bandwidth-remaining ratio. The router determines the minimum committed information rate (CIR) based on the configuration.

Examples

The following example shows how to configure a bandwidth-remaining ratio on an ATM subinterface. In the example, the router guarantees a peak cell rate of 50 Mbps for the variable bit rate-non-real time (VBR-nrt) PVC 0/200. During periods of congestion, the subinterface receives a share of excess bandwidth (unused by priority traffic) based on the bandwidth-remaining ratio of 10, relative to the other subinterfaces configured on the physical interface.

```
policy-map Child
  class precedence_0
    bandwidth 10000
  class precedence_1
    shape average 100000
    bandwidth 100
!
policy-map Parent
  class class-default
    bandwidth remaining ratio 10
    shape average 20000000
    service-policy Child
!
interface ATM2/0/3.200 point-to-point
 ip address 10.20.1.1 255.255.255.0
 pvc 0/200
  protocol ip 10.20.1.2
  vbr-nrt 50000
  encapsulation aal5snap
  service-policy output Parent
```

The following example shows how to configure bandwidth remaining ratios for individual class queues. Some of the classes configured have bandwidth guarantees and a bandwidth-remaining ratio explicitly specified. When congestion occurs within a subinterface level, the class queues receive excess bandwidth (unused by priority traffic) based on their class-level bandwidth-remaining ratios: 20, 30, 120, and 100, respectively for the precedence_0, precedence_1, precedence_2, and precedence_5 classes. Normally, the precedence_3 class (without a defined ratio) would receive bandwidth based on the bandwidth-remaining ratio of the class-default class defined in the Child policy. However, in the example, the Child policy does not define a class-default bandwidth remaining ratio, therefore, the router uses a ratio of 1 to allocate excess bandwidth to precedence_3 traffic.

```

policy-map Child
  class precedence_0
    shape average 100000
    bandwidth remaining ratio 20
  class precedence_1
    shape 10000
    bandwidth remaining ratio 30
  class precedence_2
    shape average 200000
    bandwidth remaining ratio 120
  class precedence_3
    set ip precedence 3
  class precedence_5
    set ip precedence 5
    bandwidth remaining ratio 100
policy-map Parent
  class class-default
    bandwidth remaining ratio 10
    service-policy Child
!
interface GigabitEthernet 2/0/1.10
  encapsulation dot1q 10
  service-policy output Parent

```

Related Commands

Command	Description
show policy-map	Displays the configuration of all classes for a specified service policy map or all classes for all existing policy maps. If configured, the command output includes information about ATM overhead accounting and bandwidth-remaining ratios, used to determine a queue's fair share of excess bandwidth during congestion.
show policy-map interface	Displays the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface or on a specific PVC on the interface. If configured, the command output includes information about bandwidth-remaining ratios, used to determine a queue's fair share of excess bandwidth during congestion.

show policy-map

To display the configuration of all classes for a specified service policy map or all classes for all existing policy maps, use the **show policy-map** command in EXEC mode.

show policy-map [*policy-map*]

Syntax Description	<i>policy-map</i>	(Optional) Name of the service policy map whose complete configuration is to be displayed.
---------------------------	-------------------	--

Command Default	All existing policy map configurations are displayed.
------------------------	---

Command Modes	EXEC
----------------------	------

Command History	Release	Modification
	12.0(5)T	This command was introduced.
	12.0(5)XE	This command was integrated into Cisco IOS Release 12.0(5)XE.
	12.0(7)S	This command was integrated into Cisco IOS Release 12.0(7)S.
	12.1(1)E	This command was integrated into Cisco IOS Release 12.1(1)E.
	12.2(13)T	The output of this command was modified for the Percentage-Based Policing and Shaping feature and includes the bandwidth percentage used when calculating traffic policing and shaping.
	12.0(28)S	The output of this command was modified for the QoS: Percentage-Based Policing feature to display the committed (conform) burst (Bc) and excess (peak) burst (Be) sizes in milliseconds (ms).
	12.2(14)SX	Support for this command was introduced on the Supervisor Engine 720.
	12.2(17d)SXB	Support for this command on the Supervisor Engine 2 was extended to Cisco IOS Release 12.2 SX.
	12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
	12.2(31)SB2	This command was enhanced to display bandwidth-remaining ratios configured on traffic classes and ATM overhead accounting, and was implemented on the Cisco 10000 series router for the PRE3.

Usage Guidelines	The show policy-map command displays the configuration of a service policy map created using the policy-map command. You can use the show policy-map command to display all class configurations comprising any existing service policy map, whether or not that service policy map has been attached to an interface. The command output includes bandwidth-remaining ratio configuration and statistical information, if configured and used to determine the amount of unused (excess) bandwidth to allocate to a class queue during periods of congestion.
-------------------------	---

Examples

The following is sample output from the **show policy-map** command. This sample output displays the contents of a policy map called “policy1.” In policy 1, traffic policing on the basis of a committed information rate (CIR) of 20 percent has been configured, and the bc and be have been specified in milliseconds. As part of the traffic policing configuration, optional conform, exceed, and violate actions have been specified.

```
Router# show policy-map policy1

Policy Map policy1
Class class1
  police cir percent 20 bc 300 ms pir percent 40 be 400 ms
    conform-action transmit
    exceed-action drop
    violate-action drop
```

Table 5 describes the significant fields shown in the display.

Table 5 *show policy-map Field Descriptions*

Field	Description
Policy Map	Name of policy map displayed.
Class	Name of the class configured in the policy map displayed.
police	Indicates that traffic policing on the basis of specified percentage of bandwidth has been enabled. The committed burst (Bc) and excess burst (Be) sizes have been specified in milliseconds (ms), and optional conform, exceed, and violate actions have been specified.

Bandwidth-Remaining Ratio Example

The following sample output for the **show policy-map** command indicates that the class-default class of the policy map named vlan10_policy has a bandwidth-remaining ratio of 10. When congestion occurs, the scheduler allocates class-default traffic 10 times the unused bandwidth allocated in relation to other subinterfaces.

```
Router# show policy-map vlan10_policy
Policy Map vlan10_policy
Class class-default
  Average Rate Traffic Shaping
  cir 1000000 (bps)
  bandwidth remaining ratio 10
  service-policy child_policy
```

ATM Overhead Accounting Example

The following sample output for the **show policy-map** command indicates that ATM overhead accounting is enabled for the class-default class. The BRAS-DSLAM encapsulation is dot1q and the subscriber encapsulation is snap-rbe for the AAL5 service.

```
Policy Map unit-test
Class class-default
  Average Rate Traffic Shaping
  cir 10% account dot1q aal5 snap-rbe
```

Table 6 describes the significant fields shown in the display.

Table 6 *show policy-map Field Descriptions for ATM Overhead Accounting*

Field	Description
Average Rate	Committed burst (Bc) is the maximum number of bits sent out in each interval.
cir 10%	Committed information rate (CIR) is 10 percent of the available interface bandwidth.
dot1q	BRAS-DSLAM encapsulation is 802.1Q VLAN.
aal5	DSLAM-CPE encapsulation type is based on the ATM Adaptation Layer 5 service. AAL5 supports connection-oriented variable bit rate (VBR) services.
snap-rbe	Subscriber encapsulation type.

Related Commands

Command	Description
bandwidth	Specifies or modifies the bandwidth allocated for a class belonging to a policy map, and enables ATM overhead accounting.
bandwidth remaining ratio	Specifies a bandwidth-remaining ratio for class queues and subinterface-level queues to determine the amount of unused (excess) bandwidth to allocate to the queue during congestion.
class (policy map)	Specifies the name of the class whose policy you want to create or change, and the default class (commonly known as the class-default class) before you configure its policy.
class-map	Creates a class map to be used for matching packets to a specified class.
policy-map	Creates or modifies a policy map that can be attached to one or more interfaces to specify a service policy.
shape	Shapes traffic to the indicated bit rate according to the algorithm specified, and enables ATM overhead accounting.
show policy-map interface	Displays the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface or on a specific PVC on the interface. If configured, the command output includes information about bandwidth-remaining ratios, used to determine a queue's fair share of excess bandwidth during congestion.
show running-config	Displays the current configuration of the router. If configured, the command output includes information about ATM overhead accounting.

show policy-map interface

To display the packet statistics of all classes and all priority levels configured for all service policies either on the specified interface or subinterface or on a specific permanent virtual circuit (PVC) on the interface, use the **show policy-map interface** command in privileged EXEC mode.

```
show policy-map interface [type access-control] interface-name [vc [vpi] vci] [dlci dlci]
[input | output]
```

ATM Shared Port Adapter

```
show policy-map interface atm slot/subslot/port [.subinterface]
```

Syntax Description		
type access-control	(Optional) Displays class maps configured to determine the exact pattern to look for in the protocol stack of interest.	
<i>interface-name</i>	Name of the interface or subinterface whose policy configuration is to be displayed.	
vc	(Optional) For ATM interfaces only, shows the policy configuration for a specified PVC. The name can be up to 16 characters long.	
<i>vpi</i>	(Optional) ATM network virtual path identifier (VPI) for this PVC. On the Cisco 7200 and 7500 series routers, this value ranges from 0 to 255. The <i>vpi</i> and <i>vci</i> arguments cannot both be set to 0; if one is 0, the other cannot be 0.	
<i>vci</i>	(Optional) ATM network virtual channel identifier (VCI) for this PVC. This value ranges from 0 to 1 less than the maximum value set for this interface by the atm vc-per-vp command. Typically, the lower values 0 to 31 are reserved for specific traffic (F4 Operation, Administration, and Maintenance (OAM), switched virtual circuit (SVC) signaling, Integrated Local Management Interface (ILMI), and so on) and should not be used. The VCI is a 16-bit field in the header of the ATM cell. The VCI value is unique only on a single link, not throughout the ATM network, because it has local significance only. The <i>vpi</i> and <i>vci</i> arguments cannot both be set to 0; if one is 0, the other cannot be 0.	
dlci	(Optional) Indicates that a specific PVC for which policy configuration will be displayed.	
<i>dlci</i>	(Optional) A specific data-link connection identifier (DLCI) number used on the interface. Policy configuration for the corresponding PVC will be displayed when a DLCI is specified.	
input	(Optional) Indicates that the statistics for the attached input policy will be displayed.	
output	(Optional) Indicates that the statistics for the attached output policy will be displayed.	

<i>slot</i>	(ATM Shared Port Adapter only) Chassis slot number. Refer to the appropriate hardware manual for slot information. For SIPs, refer to the platform-specific SPA hardware installation guide or the corresponding “Identifying Slots and Subslots for SIPs, SSCs, and SPAs” topic in the platform-specific SPA software configuration guide.
<i>/subslot</i>	(ATM Shared Port Adapter only) Secondary slot number on a SPA interface processor (SIP) where a SPA is installed. Refer to the platform-specific SPA hardware installation guide and the corresponding “Specifying the Interface Address on a SPA” topics in the platform-specific SPA software configuration guide for subslot information.
<i>/port</i>	(ATM Shared Port Adapter only) Port or interface number. Refer to the appropriate hardware manual for port information. For SPAs, refer to the corresponding “Specifying the Interface Address on a SPA” topics in the platform-specific SPA software configuration guide.
<i>.subinterface</i>	(ATM Shared Port Adapter only—Optional) Subinterface number. The number that precedes the period must match the number to which this subinterface belongs. The range is 1 to 4,294,967,293.

Defaults

The absence of both the forward slash (/) and a *vpi* value defaults the *vpi* value to 0. If this value is omitted, information for all virtual circuits (VCs) on the specified ATM interface or subinterface is displayed.

ATM Shared Port Adapter

When used with the ATM shared port adapter, this command has no default behavior or values.

Command Modes

Privileged EXEC

ATM Shared Port Adapter

When used with the ATM shared port adapter, EXEC or privileged EXEC.

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.0(5)XE	This command was integrated into Cisco IOS Release 12.0(5)XE.
12.0(7)S	This command was integrated into Cisco IOS Release 12.0(7)S.
12.1(1)E	This command was integrated into Cisco IOS Release 12.1(1)E.
12.1(2)T	This command was modified to display information about the policy for all Frame Relay PVCs on the interface, or, if a DLCI is specified, the policy for that specific PVC. This command was also modified to display the total number of packets marked by the quality of service (QoS) set action.
12.1(3)T	This command was modified to display per-class accounting statistics.
12.2(4)T	This command was modified for two-rate traffic policing. It now can display burst parameters and associated actions.

Release	Modification
12.2(8)T	<p>The command was modified for the Policer Enhancement—Multiple Actions feature and the WRED—Explicit Congestion Notification (ECN) feature.</p> <p>For the Policer Enhancement—Multiple Actions feature, the command was modified to display the multiple actions configured for packets conforming to, exceeding, or violating a specific rate.</p> <p>For the WRED—Explicit Congestion Notification (ECN) feature, the command displays ECN marking information</p>
12.2(13)T	<p>The following modifications were made:</p> <ul style="list-style-type: none"> • This command was modified for the Percentage-Based Policing and Shaping feature. • This command was modified for the Class-Based RTP and TCP Header Compression feature. • This command was modified as part of the Modular QoS CLI (MQC) Unconditional Packet Discard feature. Traffic classes in policy maps can now be configured to discard packets belonging to a specified class. • This command was modified to display the Frame Relay DLCI number as a criterion for matching traffic inside a class map. • This command was modified to display Layer 3 packet length as a criterion for matching traffic inside a class map. • This command was modified for the Enhanced Packet Marking feature. A mapping table (table map) can now be used to convert and propagate packet-marking values.
12.2(15)T	This command was modified to display Frame Relay voice-adaptive traffic-shaping information.
12.0(28)S	This command was modified for the QoS: Percentage-Based Policing feature to include milliseconds when calculating the committed (conform) burst (bc) and excess (peak) burst (be) sizes.
12.3(14)T	This command was modified to display bandwidth estimation parameters.
12.2(18)SXE	This command was integrated into Cisco IOS Release 12.2(18)SXE. This command was modified to display aggregate WRED statistics for the ATM shared port adapter. Note that changes were made to the syntax, defaults, and command modes. These changes are labelled “ATM Shared Port Adapter” in this document.
12.4(4)T	The type access-control keywords were added to support flexible packet matching.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and its output was modified to display either legacy (nondistributed processing) QoS or hierarchical queueing framework (HQF) parameters on Frame Relay interfaces or PVCs.
12.2(31)SB2	This command was enhanced to display statistical information for each level of priority service configured and information about bandwidth-remaining ratios, and was implemented on the Cisco 10000 series router for the PRE3.

Usage Guidelines

The **show policy-map interface** command displays the packet statistics for classes and priority levels on the specified interface or the specified PVC only if a service policy has been attached to the interface or the PVC. The command output includes bandwidth-remaining ratios configured on traffic classes.

You can use the *interface-name* argument to display output for a PVC only for enhanced ATM port adapters (for example, the PA-A3) that support per-VC queueing.

The counters displayed after the **show policy-map interface** command is entered are updated only if congestion is present on the interface.

The **show policy-map interface** command displays policy information about Frame Relay PVCs only if Frame Relay Traffic Shaping (FRTS) is enabled on the interface.

The **show policy-map interface** command displays ECN marking information only if ECN is enabled on the interface.

To determine if shaping is active with the hierarchical queuing framework (HQF), check the queue depth field of the "(queue depth/total drops/no-buffer drops)" line in the **show policy-map interface** command output.

Examples

Example of Multiple Priority Queues on Serial Interface

The following sample output from the **show policy-map interface** command shows the types of statistical information that displays when multiple priority queues are configured. Depending upon the interface in use and the options enabled, the output you see may vary slightly from the output shown below.

```
Router# show policy-map interface

Serial2/1/0
Service-policy output: P1
Queue statistics for all priority classes:
.
.
.
Class-map: Gold (match-all)
  0 packets, 0 bytes/*Updated for each priority level configured.*/
  5 minute offered rate 0 bps, drop rate 0 bps
Match: ip precedence 2
  Priority: 0 kbps, burst bytes 1500, b/w exceed drops: 0
Priority Level 4:
  0 packets, 0 bytes
```

Example of Bandwidth-Remaining Ratios

The following sample output from the **show policy-map interface** command indicates that bandwidth-remaining ratios are configured for class queues. As shown in the example, the classes precedence_0, precedence_1, and precedence_2 have bandwidth-remaining ratios of 20, 40, and 60, respectively.

```
Router# show policy-map interface GigabitEthernet1/0/0.10

Service-policy output: vlan10_policy

Class-map: class-default (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
Match: any
  0 packets, 0 bytes
  30 second rate 0 bps
Queueing
```

```

queue limit 250 packets
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 0/0
shape (average) cir 1000000, bc 4000, be 4000
target shape rate 1000000
bandwidth remaining ratio 10

Service-policy : child_policy

Class-map: precedence_0 (match-all)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
  Match: ip precedence 0
  Queueing
    queue limit 62 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 0/0
    shape (average) cir 500000, bc 2000, be 2000
    target shape rate 500000
    bandwidth remaining ratio 20

Class-map: precedence_1 (match-all)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
  Match: ip precedence 1
  Queueing
    queue limit 62 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 0/0
    shape (average) cir 500000, bc 2000, be 2000
    target shape rate 500000
    bandwidth remaining ratio 40

Class-map: precedence_2 (match-all)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
  Match: ip precedence 2
  Queueing
    queue limit 62 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 0/0
    shape (average) cir 500000, bc 2000, be 2000
    target shape rate 500000
    bandwidth remaining ratio 60

Class-map: class-default (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps, drop rate 0 bps
  Match: any
    0 packets, 0 bytes
    30 second rate 0 bps

queue limit 62 packets
(queue depth/total drops/no-buffer drops) 0/0/0
(pkts output/bytes output) 0/0

```

Related Commands	Command	Description
	bandwidth remaining ratio	Specifies a bandwidth-remaining ratio for class queues and subinterface-level queues to determine the amount of unused (excess) bandwidth to allocate to the queue during congestion.
	priority	Specifies that low-latency behavior must be given to a traffic class and configures multiple priority queues.
	police	Configures traffic policing.
	police (percent)	Configures traffic policing on the basis of a percentage of bandwidth available on an interface.
	police (two rates)	Configures traffic policing using two rates, the committed information rate (CIR) and the peak information rate (PIR).
	policy-map	Creates or modifies a policy map that can be attached to one or more interfaces to specify a service policy.
	show interfaces	Displays statistics for all interfaces configured on a router or access server.
	show policy-map	Displays the configuration of all classes for a specified service policy map or all classes for all existing policy maps. If configured, the command output includes information about ATM overhead accounting and bandwidth-remaining ratios, used to determine a queue's fair share of excess bandwidth during congestion.
	show policy-map class	Displays the configuration for the specified class of the specified policy map.

CCVP, the Cisco Logo, and the Cisco Square Bridge logo 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, BPX, 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, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, *Packet*, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, 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. (0609R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2006 Cisco Systems, Inc. All rights reserved.

■ show policy-map interface