Distributing Bandwidth Between Queues

Developing a queuing strategy is an important step in optimizing network functionality and services. Equally important is ensuring that bandwidth is shared fairly among the competing traffic queues. By distributing bandwidth between the different queues of a link, you can increase the capacity of a link and optimize aspects of your network that have greater importance.

The Cisco 10000 series router allows you to distribute bandwidth to ensure that bandwidth is shared fairly between the queues of a link. This chapter describes bandwidth distribution between queues on the Cisco 10000 series router and contains the following topics:

Bandwidth Distribution Between Queues

Distributing bandwidth on a link using the **bandwidth** command ensures that bandwidth is shared fairly among competing traffic. The router uses class queues to allocate bandwidth, first servicing priority queue traffic followed by either bandwidth guarantee or bandwidth remaining queue traffic. By default,

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

a minimum bandwidth guaranteed queue has buffers for up to 50 milliseconds of 256-byte packets at line rate, but not less than 32 packets. The router does not ensure latency characteristics for bandwidth queues.

The router can commit up to 99 percent of the interface bandwidth to one or more class queues. If you attempt to attach a policy map to an interface when the sum of the bandwidth assigned to classes is greater than 99 percent of the available bandwidth, the router logs a warning message and does not allocate the requested bandwidth to all of the classes. If the policy map is already attached to other interfaces, it is removed from them.

Bandwidth includes the Layer 2 header, Layer 2 payload, and two bytes of trailer.

- On ATM networks, the bandwidth is cell-based and includes Layer 2 overhead and cell overhead (for example, cell overhead for SNAP and AAL5, the cell header, the AAL5 trailer, and AAL5 padding).
- On Frame Relay networks with link fragmentation and interleaving (LFI) enabled, bandwidth is based on fragments with Layer 2 overhead included.

The router converts the specified bandwidth to the nearest multiple of 1/255 (PRE1) or 1/65535 (PRE2) of the interface speed. Use the **show policy-map interface** command to display the actual bandwidth.

Feature History for Bandwidth Distribution Between Queues

Cisco IOS Release	Description	Required PRE
Release 12.0(17)SL	This feature was introduced on the PRE1.	PRE1
Release 12.2(15)BX	This feature was introduced on the PRE2.	PRE2
Release 12.2(28)SB	This feature was integrated in Cisco IOS Release 12.2(28)SB for the PRE2.	PRE2
Release 12.2(31)SB2	This feature was introduced on the PRE3. The Distribution of Remaining Bandwidth Using Ratio feature was also introduced on the PRE3 only.	PRE2 PRE3

Interfaces Supporting Bandwidth Distribution

The following describes interface support for bandwidth distribution using the **bandwidth** command:

Interfaces Supporting the bandwidth Command (Outbound Only)

- Physical
- Multilink PPP and Multilink Frame Relay
- ATM shaped (peak cell rate is specified) unspecified bit rate (UBR) PVCs and point-to-point subinterfaces
- ATM constant bit rate (CBR) PVCs and point-to-point subinterfaces
- ATM variable bit rate (VBR) PVCs and point-to-point subinterfaces
- Label-controlled ATM (LC-ATM) subinterfaces *
- Frame Relay PVCs, point-to-point subinterfaces, and map classes *
- Ethernet VLANs *

* Requires a specific type of hierarchical policy. For more information, see the Chapter 12, "Defining QoS for Multiple Policy Levels."

Note The router only supports the **bandwidth** command on outbound interfaces.

Interfaces Not Supporting the bandwidth Command

- ATM unshaped (no peak cell rate specified) UBR PVCs and point-to-point subinterfaces
- IP tunnel
- Virtual-access (See the "VAI QoS Inheritance" section on page 4-24.)

Note The router does not support the **bandwidth** command on inbound interfaces.

Unused Bandwidth Allocation

After the Cisco 10000 series router allocates bandwidth to priority and bandwidth guaranteed class queues, the router divides unused (excess) bandwidth among the packets remaining in the class queues.

- For all releases prior to Cisco IOS Release 12.0(19)SL, the router divides the unused bandwidth equally among the class queues with outstanding packets.
- For Cisco IOS Release 12.0(19)SL and later releases, the router divides unused bandwidth proportional to the class bandwidth guarantee. You can override this proportional distribution by using the **bandwidth remaining percent** command. The router distributes unused bandwidth only to non-priority queues; a priority queue never receives more than its guaranteed bandwidth.

In Example 22, the policy-map named VLAN guarantees 30 percent of the bandwidth to the class named Customer1 and 60 percent of the bandwidth to the class named Customer2. If you apply the VLAN policy map to a 1-Mbps link, 300 kbps is guaranteed to class Customer1 and 600 kbps is guaranteed to class Customer2, with 100 kbps remaining for the class-default class. If the class-default class does not need additional bandwidth, the unused 100 kbps is available for use by class Customer1 and class Customer2. If both classes need the bandwidth, they share it in proportion to the configured rates. In this example, the sharing ratio is 30:60 or 1:2.

Example 22 Excess Bandwidth Allocation

```
Router(config)# policy-map VLAN
Router(config-pmap)# class Customer1
Router(config-pmap-c)# bandwidth percent 30
Router(config-pmap-c)# exit
Router(config-pmap)# class Customer2
Router(config-pmap-c)# bandwidth percent 60
```

Traffic Classes That Can Use Excess Bandwidth

The **bandwidth** and **priority** commands provide bandwidth guarantees that are often described as bandwidth that is reserved or set aside. However, neither command implements a true reservation of bandwidth. If a traffic class is not using its configured bandwidth, the unused bandwidth is shared among the other classes.

The Cisco 10000 series router guarantees a priority class 95 percent of the bandwidth. As a result, the priority class can starve other traffic classes unless you configure policing. The router does not police the priority class unless you configure a policer.

Table 1 describes when a class configured with the **bandwidth** or **priority** command can use excess bandwidth.

Table 1 Comparing When a Bandwidth and Priority Class Can Use Excess Bandwidth

Command	Congestion	Non-Congestion
bandwidth	Allowed to exceed the allocated rate.	Allowed to exceed the allocated rate.
priority	Not allowed to exceed the configured rate. If you configure a policer for the priority class, the router polices the priority class to the configured bps or percentage rate, and discards any excess packets. If you do not configure a policer, the router guarantees a priority class 95 percent of the bandwidth, which can result in bandwidth starvation of other traffic classes.	Not allowed to exceed the configured rate. If you configure a policer, the router polices the priority class to the configured bps or percentage rate, and discards any excess packets.

Bandwidth Calculations

The Cisco 10000 series router can commit up to 99 percent of the interface bandwidth to one or more class queues. If you attempt to attach a policy map to an interface when the sum of the bandwidth assigned to classes is greater than 99 percent of the available bandwidth, the router logs a warning message and does not allocate the requested bandwidth to all of the classes. If the policy map is already attached to other interfaces, it is removed from them.

For a physical interface, the total bandwidth is the bandwidth of the physical interface. The router converts the minimum bandwidth that you specify to the nearest multiple of 1/255 (PRE1) or 1/65535 (PRE2) of the interface speed. When you request a value that is not a multiple of 1/255 or 1/65535, the router chooses the nearest multiple.

The bandwidth percentage is based on the interface bandwidth. In a hierarchical policy, the bandwidth percentage is based on the nearest parent shape rate.

By default, a minimum bandwidth guaranteed queue has buffers for up to 50 milliseconds of 256-byte packets at line rate, but not less than 32 packets.

Table 2 describes the data included and excluded in the committed rate when a traffic shaper and policer are configured on the interface.

Table 2 Committed Rate Data for Policing and Shaping

Action	Frame Relay	**Ethernet**	**ATM**	**ADSL**	**SDSL**
Policing	No bit or byte stuffing No 7E flags[1] No Frame Check Sequence (FCS)	**Generic Traffic Shaping** **Class-Based Shaping** No Inter-Frame Gap (IFG) No Preamble No Start of Frame Delimiter (SFD) No Frame Check Sequence (FCS)	No cell header No AAL Common Part Convergence Sublayer (CPCS) pad No ATM trailer	**Class-Based Shaping** No ATM cell overhead No AAL Common Part Convergence Sublayer (CPCS) pad No ATM trailer	**Class-Based Shaping** No ATM cell overhead No AAL Common Part Convergence Sublayer (CPCS) pad No ATM trailer
Shaping	4 bytes of Frame Check Sequence (FCS) No bit or byte stuffing No 7E flags	**Generic Traffic Shaping** **Class-Based Shaping** No Inter-Frame Gap (IFG) No Preamble No Start of Frame Delimiter (SFD) Frame Check Sequence (FCS)	ATM cell header AAL Common Part Convergence Sublayer (CPCS) pad ATM trailer	**Class-Based Shaping** No ATM cell overhead No AAL Common Part Convergence Sublayer (CPCS) pad	**Class-Based Shaping** No ATM cell overhead No AAL Common Part Convergence Sublayer (CPCS) pad

1. The router does not account for flags or Frame Check Sequence (FCS) that the hardware adds or removes.

Table 3 describes what bandwidth is based on for each media type.

Table 3 Basis for Bandwidth

Media	**Bandwidth Based On**
Frame Relay	Fragments[1]
Ethernet	Bits
ATM variable bit rate (VBR)	Sustained cell rate (SCR)
ATM unspecified bit rate (UBR)	Peak cell rate (PCR)
ATM constant bit rate (CBR)	Peak cell rate (PCR)

1. For Frame Relay networks with link fragmentation and interleaving (LFI) enabled.

Bandwidth Allocation to PVCs

The router allocates bandwidth to variable bit rate (VBR), constant bit rate (CBR), and shaped (specifies the peak cell rate) unspecified bit rate (UBR) PVCs before allocating bandwidth to unshaped UBR PVCs. As a result, a diminished amount of bandwidth is available to allocate to unshaped UBR PVCs. To override this behavior, create a hierarchical (nested) policy with the bandwidth specified and attach the policy to the ATM port or physical interface.

Class-Default Bandwidth Guarantee

The Cisco 10000 series router guarantees bandwidth for the class-default class in the following way:

- For all releases prior to Cisco IOS Release 12.0(19)SL, when other classes are present, the class-default class receives no bandwidth guarantee unless it has a **bandwidth** command configured.
- Cisco IOS Release 12.0(19)SL adds support for an implicit bandwidth guarantee. In Release 12.0(19)SL and later releases, the class-default class receives a bandwidth guarantee of any uncommitted interface bandwidth plus one percent. You can decrease this guarantee by using the **bandwidth** command in the class-default class.

Note You can decrease the class-default bandwidth guarantee, but you cannot increase it.

Committed Rate Data

The committed rate includes the framing overhead, excluding bit and byte stuffing. For ATM, the committed rate also includes the cell overhead. This causes the actual packet data rate to be lower than what you specify in the **bandwidth** command. Therefore, be sure to allow for the framing and cell overhead when you specify a minimum bandwidth for a class.

For example, if you need to commit a rate of 1000 64-byte packets per second and each have four bytes of framing overhead, instead of specifying 512 kbps in the **bandwidth** command, you would specify 544 kbps, which is calculated in the following way:

```
1000 * (64+4) * 8 / 1000 = 544
```

A similar setup on ATM would require 848 kbps because each 64-byte packet requires two cells of 53 bytes:

```
1000 * 2 * 53 * 8 / 1000 = 848
```

Bandwidth and Priority Service

To configure priority queuing, use the **priority** command in policy-map class configuration mode and configure the bandwidth in the following way:

- In releases prior to Cisco IOS Release 12.3(7)XI, specify a minimum bandwidth rate using the **priority** command.
- In Cisco IOS Release 12.3(7)XI and later releases, specify a minimum bandwidth rate using the **police** command. The router no longer supports the *percent* and *bandwidth-kbps* options for the **priority** command.

When you create a priority queue using the **priority** command, the router sets the committed information rate (CIR) of the queue to near link bandwidth. Therefore, the priority queue can consume almost all of the link bandwidth when packets are transmitted from it. As a result, there is no way to guarantee bandwidth to the other queues on the link. To resolve this, configure the other queues using the **bandwidth remaining** command. This command enables the router to allocate relative amounts of bandwidth, eliminating the need to constantly change the class bandwidth.

Table 4 lists the functional differences between the **bandwidth** and **priority** commands:

Table 4 Comparison of Functional Differences Between the bandwidth and priority Commands

Function	bandwidth Command	priority Command
Minimum bandwidth guarantee	Yes	Yes
Maximum bandwidth guarantee	No	Yes
Built-in policer	No	No
Provides low-latency	No	Yes

The **bandwidth** and **priority** commands are also designed to meet different QoS policy objectives. Table 5 lists those differing objectives:

Table 5 Comparison of QoS Policy Objectives of the bandwidth and priority Commands

Application	bandwidth Command	priority Command
Bandwidth management for WAN links	Yes	Somewhat
Manage delay and variations in delay (jitter)	No	Yes
Improve application response time	No	Yes

Queuing Modes

The Cisco 10000 series router supports the following queuing modes:

- atm pxf queuing Mode, page 7
- no atm pxf queuing Mode, page 8

atm pxf queuing Mode

The atm pxf queuing mode supports low VC counts. Before you configure VCs on an interface, configure the queuing mode for the port (atm pxf queuing or no atm pxf queuing). After you configure the mode, then configure the VCs. Do not change the mode while VCs are configured on the interface. If you need to change the mode, delete the VCs first and then change the mode. Changing the mode while VCs are configured can produce undesired results.

To support a low VC count, configure the **atm pxf queuing** command on each port of the Cisco 10000 series router. For Cisco IOS Release 12.3(7)XI2 and later releases, all line cards support a maximum of 28,672 VCs when configured for hierarchical shaping.

Note For releases prior to Cisco IOS Release 12.3(7)XI2, the OC-3 and OC-12 line cards support a maximum of 14,336 VCs when configured for hierarchical shaping. The DS3/E3 line card supports a maximum of 8,192 VCs when configured for shaping. You can configure the maximum number of VCs across the ports in any fashion, provided that you do not exceed the per-port maximum. The OC-3 line card is limited to 8,192 VCs per port and the DS3 is limited to 4,096 VCs per port.

The Cisco 10000 series router supports the following ATM traffic classes when you configure the **atm pxf queuing command**:

- Shaped and unshaped unspecified bit rate (UBR)

To configure shaped UBR, enter the **ubr** command and specify the PCR value. For unshaped UBR, do not specify the PCR value.

- Variable bit rate-nonreal time (VBR-nrt)

  To configure VBR-nrt QoS, use the **vbr-nrt** command and specify the output PCR, output sustainable cell rate (SCR), and output maximum burst cell size (MBS).

- Constant bit rate (CBR)

  To configure CBR QoS, use the **cbr** command and specify the average cell rate for the ATM circuit emulation service (CES) for an ATM permanent virtual circuit (PVC).

Note If you use the **ubr+** command to configure shaped UBR, the router accepts the PCR value you specify, but it does not use it. The router does not notify you of this behavior.

no atm pxf queuing Mode

Note Cisco recommends that you do not configure no atm pxf queuing mode for QoS-sensitive deployments.

The no atm pxf queuing mode supports high VC counts. Before you configure VCs on an interface, configure the queuing mode for the port (atm pxf queuing or no atm pxf queuing). After you configure the mode, then configure the VCs. Do not change the mode while VCs are configured on the interface. If you need to change the mode, delete the VCs first and then change the mode. Changing the mode while VCs are configured can produce undesired results.

To support a high number of virtual circuits (VCs), configure the **no atm pxf queuing** command on each port of the Cisco 10000 series router. PPPoA supports one session per VC and requires that you enable no atm pxf queuing mode to support up to 64,000 PPPoA sessions. Enabling no atm pxf queuing mode is not required for L2TP and might not be required for PPPoE because you can have 64,000 sessions on a single VC.

The Cisco 10000 series router supports the following ATM traffic classes when you configure the **no atm pxf queuing command**:

- Shaped and unshaped unspecified bit rate (UBR)

  To configure shaped UBR, enter the **ubr** command and specify the PCR value. For unshaped UBR, enter the **ubr** command without specifying the PCR value.

- Variable bit rate-nonreal time (VBR-nrt)

  To configure VBR-nrt QoS, use the **vbr-nrt** command and specify the output PCR, output sustainable cell rate (SCR), and the output maximum burst cell size (MBS) for a VC class.

Note If you use the **ubr+** command to configure shaped UBR, the router accepts the PCR value you specify, but it does not use it. The router does not notify you of this behavior.

Restrictions and Limitations for Bandwidth Distribution

- You cannot configure the **bandwidth** command on a class with priority service configured.
- In a policy map, you can configure the **bandwidth**, **priority**, or **shape** command for a maximum of 14 (PRE1) or 30 (PRE2) non-class-default classes.

- The router supports a maximum of 32,767 (PRE1) or 128,000 (PRE2) output packet queues.
- A policy map can have only one priority queue.

Configuring Bandwidth Distribution Between Queues

To configure or modify the amount of bandwidth allocated for a traffic class, enter the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)#` **`policy-map`** *`policy-map-name`*	Specifies the name of the policy map and enters policy-map configuration mode. *policy-map-name* is the name of the policy map.
Step 2	`Router(config-pmap)#` **`class`** *`class-map-name`*	Assigns the traffic class you specify to the policy map. Enters policy-map class configuration mode. *class-map-name* is the name of a previously configured class map and is the traffic class for which you want to define QoS actions.
Step 3	`Router(config-pmap-c)#` **`bandwidth`** {*`bandwidth-kbps`* | **`percent`** *`percentage`* | ***`remaining percent`*** *`percentage`*}	Specifies or modifies the minimum bandwidth allocated for a traffic class in a policy map. *bandwidth-kbps* specifies or modifies the minimum bandwidth allocated for a class belonging to a policy map. Valid values are from 8 to 2,488,320, which represents from 1 to 99 percent of the link bandwidth. **Note** The range of valid values for *bandwidth-kbps* might be smaller than the values indicated above. Use the question mark (?) in context-sensitive help to display the range of valid values. **percent** *percentage* specifies or modifies the minimum percentage of the link bandwidth allocated for a class belonging to a policy map. Valid values are from 1 to 99. **remaining percent** *percentage* specifies or modifies the minimum percentage of unused link bandwidth allocated for a class belonging to a policy map. Valid values are from 1 to 99. **Note** Configure the amount of bandwidth large enough to also accommodate Layer 2 overhead.

For information about classifying traffic and creating QoS service policies, see Chapter 2, “Classifying Traffic” and Chapter 1, “Configuring QoS Policy Actions and Rules.”

Configuration Examples for Configuring Bandwidth Distribution

This section provides the following configuration examples:

- Configuration Example for Kilobits per Second-Based Bandwidth, page 10

Configuration Example for Kilobits per Second-Based Bandwidth

Example 23 shows how to create a policy-map named account1 and configure a class named gold, which provides a guaranteed minimum bandwidth of 700 kbps for all gold class traffic. The maximum queue depth for gold traffic is 64 packets. If the queue depth is reached, the router tail-drops excess traffic.

Example 23 Configuring Bandwidth Based on Kilobits per Second

```
Router(config)# policy-map account1
Router(config-pmap)# class gold
Router(config-pmap-c)# bandwidth 700
Router(config-pmap-c)# queue-limit 64
```

Configuration Example for Percentage-Based Bandwidth

Example 24 shows how to create a policy map named mypolicy and configure a class named silver, which allocates a minimum of 30 percent of the link bandwidth to all silver class traffic. The maximum queue depth for silver traffic is 64 packets. The **random-detect** command ensures that excess traffic is randomly discarded using a precedence-based algorithm.

Example 24 Configuring a Percentage-Based Bandwidth

```
Router(config)# policy-map mypolicy
Router(config-pmap)# class silver
Router(config-pmap-c)# bandwidth percent 30
Router(config-pmap-c)# queue-limit 64
Router(config-pmap-c)# random-detect precedence-based
```

Configuration Example for Bandwidth-Remaining-Based Bandwidth

Example 25 shows how to create a policy map named vlan and configure two traffic classes named priority and premium. The priority class contains both the **priority** and **police** commands. The premium class allocates a minimum of 25 percent of the unused bandwidth to this class.

Example 25 Configuring Bandwidth-Remaining-Based Bandwidth

```
Router(config)# policy-map vlan
Router(config-pmap)# class priority
Router(config-pmap-c)# priority
Router(config-pmap-c)# police percent 50
Router(config-pmap-c)# exit
Router(config-pmap)# class premium
Router(config-pmap-c) bandwidth remaining percent 25
```

Verifying and Monitoring Bandwidth Distribution

The Cisco 10000 series router collects statistical information about the current queue length, and the number of packets and bytes transmitted and dropped (64 bits), which you can display using the **show** commands in this section.

To verify and monitor bandwidth distribution in a policy map, enter the following commands in privileged EXEC mode:

Command	Purpose
Router# **show policy-map interface**	Displays the configuration of all classes configured for all policy maps attached to all interfaces.
Router# **show policy-map interface** *interface* [**input** | **output**]	Displays the configuration of all classes configured for all inbound or outbound policy maps attached to the specified interface. *interface* is the name of the interface or subinterface whose policy configuration you want to display. **input** indicates to display the statistics for the attached inbound policy. **output** indicates to display the statistics for the attached outbound policy. **Note** If you do not specify **input** or **output**, the router displays information about all classes that are configured for all inbound and outbound policies attached to the interface you specified.
Router# **show policy-map** *policy-map-name*	Displays the configuration of all classes contained in the policy map you specify. *policy-map-name* is the name of the policy map whose configuration information you want to display. The name can be a maximum of 40 characters. If you do not specify a *policy-map-name*, the command displays the configuration of all policy maps configured on the router.
Router# **show policy-map** *policy-map-name* **class** *class-name*	Displays the configuration of the class you specify. The policy map you specify includes this class. *policy-map-name* is the name of the policy map that contains the class configuration you want to display. *class-name* is the name of the class whose configuration you want to display. If you do not specify *class-name*, the router displays class configuration for all classes in the policy map.

Verification Example for Bandwidth Distribution

Example 26 shows the bandwidth configuration for the class named Business and the class named Non-Business in the Gold policy map.

Example 26 Verifying Class Bandwidth

```
Router# show policy-map
  Policy Map Gold
    Class Business
      bandwidth 5000
    Class Non-Business
      bandwidth 2000
```

Example 27 shows the bandwidth that the router allocated to the traffic classes in the Gold policy map.

Example 27 Verifying Bandwidth Distribution

```
Router# show policy-map interface atm 3/0/0
ATM3/0/0

  Service-policy output: Gold

    Class-map: Business (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: ip precedence 4
      Output queue: 0/128; 0/0 packets/bytes output, 0/0 drops
      Bandwidth : 4999 kbps (Weight 3)

    Class-map: Non-Business (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: ip precedence 3  5
      Output queue: 0/64; 0/0 packets/bytes output, 0/0 drops
      Bandwidth : 2001 kbps (Weight 1)

    Class-map: class-default (match-any)
      134 packets, 2760 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      Output queue: 0/4096; 3046242/75866271 packets/bytes output, 0/0 drops
```

Distribution of Remaining Bandwidth Using Ratio

The Distribution of Remaining Bandwidth Using Ratio feature allows service providers to prioritize subscriber traffic during periods of congestion. A bandwidth-remaining ratio is used to influence how the router allocates excess bandwidth (unused by priority traffic) to a class of non-priority traffic. Instead of using only bandwidth rate, the router considers configured minimum bandwidth rates, maximum bandwidth rates, and bandwidth-remaining ratios when determining excess bandwidth allocation. A bandwidth-remaining ratio adds more flexibility in prioritizing traffic and enables you to influence excess bandwidth allocation by basing the bandwidth-remaining ratio on factors other than speed.

When bandwidth-remaining ratios are not specified, the hierarchical queuing framework (HQF) scheduler on the PRE3 does the following:

- Computes a default bandwidth-remaining ratio based on the subinterface speed—ATM interfaces
- Uses the minimum bandwidth-remaining ratio allowed (currently 1 on the PRE3)—Other interfaces such as VLANs and Frame Relay DLCIs

With bandwidth-remaining ratios, service providers have more flexibility in assigning priority to subinterfaces and queues during congestion. In addition to speed, you can base the bandwidth-remaining ratio on alternative factors, such as a service product or subscription rate. In this way, for example, you can give higher weight to subinterfaces carrying business services and lower weight to subinterfaces carrying residential services. The bandwidth-remaining ratio enables the HQF scheduler to service a subinterface with a low SCR but a high bandwidth-remaining ratio more frequently than servicing a subinterface with a high SCR but a low bandwidth-remaining ratio.

The Distribution of Remaining Bandwidth Using Ratio feature is available on outbound interfaces only.

Feature History for Bandwidth-Remaining Ratios

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

Bandwidth-Remaining Ratio

A *bandwidth-remaining ratio* is a value from 1 to 1000 that is used to determine the amount of unused (excess) bandwidth to allocate to a class queue or subinterface-level queue during congestion. The router allocates excess bandwidth relative to the other class queues and subinterface-level queues configured on the physical interface. The bandwidth-remaining ratio value does not indicate a percentage. For example, a subinterface with a bandwidth-remaining ratio of 100 receives 10 times the unused (excess) bandwidth during congestion than a subinterface with a bandwidth-remaining ratio of 10.

Without bandwidth-remaining ratios, the router allocates excess bandwidth based on the following:

- Speed of the subinterface (for example, the configured SCR)—ATM subinterfaces
- Minimum bandwidth-remaining ratio allowed (currently 1 on the PRE3)—Interface types such as VLANs and Frame Relay DLCIs

With bandwidth-remaining ratios, excess bandwidth allocation can be based on factors other than the bandwidth rate (for example, service product or subscription rate).

Prerequisites for Distribution of Remaining Bandwidth Using Ratio

You must configure traffic classes using the **class-map** command.

Restrictions for Distribution of Remaining Bandwidth Using Ratio

- Bandwidth-remaining ratios are only available on outbound interfaces.
- 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 Prec1
    class precedence_0
    bandwidth remaining ratio 10
    class precedence_2
    bandwidth 1000
```

- The **bandwidth remaining ratio** command cannot coexist with another **bandwidth** 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_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 of the parent policy.
- 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 percent 10
    bandwidth remaining ratio 10
```

Configuring Bandwidth-Remaining Ratios

You can apply bandwidth-remaining ratios to different subinterfaces and to different traffic queues within a single outbound interface or subinterface.

Use the following procedures to configure bandwidth-remaining ratios:

- Configuring and Applying Bandwidth-Remaining Ratios to Subinterfaces, page 14
- Configuring and Applying Bandwidth-Remaining Ratios to Class Queues, page 16

Configuring and Applying Bandwidth-Remaining Ratios to Subinterfaces

To configure and apply bandwidth-remaining ratios to subinterfaces, enter the following commands beginning in global configuration mode:

Note You can apply bandwidth-remaining ratios to outbound subinterfaces only.

	Command or Action	Purpose
Step 1	**policy-map** *child-policy-name*	Creates or modifies a child policy map. Enters policy-map configuration mode. *child-policy-name* is the name of the child policy map.
Step 2	**class** *class-map-name*	Configures the class map that you specify. Enters policy-map class configuration mode. *class-map-name* is the name of a previously created class map.
Step 3	**bandwidth** *bandwidth-kbps*	Specifies the bandwidth, in kbps, to be allocated to this traffic class. *bandwidth-kbps* is the bandwidth in kilobits per second (kbps).
Step 4	exit	Exits policy-map class configuration mode.
Step 5	exit	Exits policy-map configuration mode.
Step 6	**policy-map** *parent-policy-name*	Creates or modifies a parent policy map. Enters policy-map configuration mode. *parent-policy-name* is the name of the parent policy map.

	Command or Action	Purpose
Step 7	`class class-default`	Configures the class-default class. Enters policy-map class configuration mode. **Note** The router interprets any features configured under the class-default class as aggregate features on the subinterface.
Step 8	**bandwidth remaining ratio** *ratio*	Specifies the bandwidth-remaining ratio for the subinterface. *ratio* is the value used to determine the amount of unused bandwidth to allocate to each queue on the subinterface during periods of congestion. The scheduler allocates the excess bandwidth relative to other subinterfaces. Valid values are 1 to 1000. The default value is platform-dependent. The router distinguishes between interface types at the subinterface level when using default bandwidth-remaining ratios. On the Cisco 10000 series router the default ratio value is 1 for VLAN subinterfaces and Frame Relay DLCIs. For ATM subinterfaces, the router computes the default based on the subinterface speed.
Step 9	**shape** {**average** | **peak**} *cir* [*bc*] [*be*]	(Optional) Shapes the average or peak rate to the rate you specify. **average** specifies average rate shaping. **peak** specifies peak rate shaping. *cir* specifies the committed information rate (CIR), in bits per second (bps). (Optional) *bc* specifies the committed burst size, in bits. (Optional) *be* specifies the excess burst size, in bits.
Step 10	**service-policy** *child-policy-name*	Applies the child policy map you specify to the traffic class. The router applies the QoS actions specified in the child policy to the traffic class. *child-policy-name* is the name of the child policy. **Note** The service-policy command typically requires that you specify the direction of the traffic using the **input** or **output** keywords. However, when applying a child policy to a parent policy, do not specify traffic direction.
Step 11	`exit`	Exits policy-map class configuration mode.
Step 12	`exit`	Exits policy-map configuration mode.

	Command or Action	Purpose
Step 13	**interface** *type slot***/***module***/***port***.***subinterface* [**point-to-point** | **multipoint**]	Creates or modifies the interface you specify. Enters subinterface configuration mode. *type* is the interface type (for example, Gigabit Ethernet). *slot***/***module***/***port***.***subinterface* is the number of the subinterface that identifies the subinterface (for example, 1/0/0.1). (Optional) **point-to-point** indicates that the subinterface is a point-to-point subinterface. (Optional) **multipoint** indicates that the subinterface is a point-to-multipoint subinterface.
Step 14	**service-policy** {**input** | **output**} *parent-policy-name*	Applies the parent policy to the subinterface. **input** indicates to apply the service policy to inbound traffic. **output** indicates to apply the service policy to outbound traffic. *parent-policy-name* is the name of the parent policy map. The router shapes the subinterface traffic to the shaping rate specified in the parent class-default class and applies the QoS actions specified in the child policy to traffic matching the traffic classes. During periods of congestion, the router uses the bandwidth-remaining ratio specified in the parent policy map to allocate unused bandwidth on this subinterface relative to other subinterfaces.

Configuring and Applying Bandwidth-Remaining Ratios to Class Queues

To configure and apply bandwidth-remaining ratios to class queues, enter the following commands beginning in global configuration mode:

	Command or Action	Purpose
Step 1	**policy-map** *child-policy-name*	Creates or modifies a child policy map. Enters policy-map configuration mode. *child-policy-name* is the name of the child policy map.
Step 2	**class** *class-map-name*	Configures the class map that you specify. Enters policy-map class configuration mode. *class-map-name* is the name of a previously created class map.

	Command or Action	Purpose
Step 3	**shape** {**average** | **peak**} *cir* [*bc*] [*be*]	(Optional) Shapes the average or peak rate to the rate you specify. **average** specifies average rate shaping. **peak** specifies peak rate shaping. *cir* specifies the committed information rate (CIR), in bits per second (bps). (Optional) *bc* specifies the committed burst size, in bits. (Optional) *be* specifies the excess burst size, in bits.
Step 4	**bandwidth remaining ratio** *ratio*	(Optional) Specifies the bandwidth-remaining ratio for the traffic class. *ratio* is the value used to determine the amount of unused bandwidth to allocate to each queue on the subinterface during periods of congestion. The scheduler allocates the excess bandwidth relative to other subinterfaces. Valid values are 1 to 1000. The default value is platform-dependent. The router makes no distinction between interface types at the class level when using the default bandwidth-remaining ratio. On the Cisco 10000 series router the default bandwidth-remaining ratio value is 1.
Step 5	exit	Exits policy-map class configuration mode.
Step 6	exit	Exits policy-map configuration mode.
Step 7	**policy-map** *parent-policy-name*	Creates or modifies a parent policy map. Enters policy-map configuration mode. *parent-policy-name* is the name of the parent policy map.
Step 8	class class-default	Configures the class-default class. Enters policy-map class configuration mode. **Note** The router interprets any features configured under the class-default class as aggregate features on the subinterface.
Step 9	**shape** {**average** | **peak**} *cir* [*bc*] [*be*]	Shapes the average or peak rate to the rate you specify. **average** specifies average rate shaping. **peak** specifies peak rate shaping. *cir* specifies the committed information rate (CIR), in bits per second (bps). (Optional) *bc* specifies the committed burst size, in bits. (Optional) *be* specifies the excess burst size, in bits.

	Command or Action	Purpose
Step 10	**bandwidth remaining ratio** *ratio*	(Optional) Specifies the bandwidth-remaining ratio for the subinterface. *ratio* is the value used to determine the amount of unused bandwidth to allocate to each queue on the subinterface during periods of congestion. The scheduler allocates the excess bandwidth relative to other subinterfaces. Valid values are 1 to 1000. The default value is platform-dependent. The router distinguishes between interface types at the subinterface level when using default bandwidth-remaining ratios. On the Cisco 10000 series router the default ratio value is 1 for VLAN subinterfaces and Frame Relay DLCIs. For ATM subinterfaces, the router computes the default based on the subinterface speed.
Step 11	**service-policy** *child-policy-name*	Applies the child policy map you specify to the traffic class. The router applies the QoS actions specified in the child policy to the traffic class. *child-policy-name* is the name of the child policy. **Note** The service-policy command typically requires that you specify the direction of the traffic using the **input** or **output** keywords. However, when applying a child policy to a parent policy, do not specify traffic direction.
Step 12	exit	Exits policy-map class configuration mode.
Step 13	exit	Exits policy-map configuration mode.
Step 14	**interface** *type slot***/***module***/***port***.***subinterface* [**point-to-point** | **multipoint**]	Creates or modifies the interface you specify. Enters subinterface configuration mode. *type* is the interface type (for example, Gigabit Ethernet). *slot***/***module***/***port***.***subinterface* is the number of the subinterface that identifies the subinterface (for example, 1/0/0.1). (Optional) **point-to-point** indicates that the subinterface is a point-to-point subinterface. (Optional) **multipoint** indicates that the subinterface is a point-to-multipoint subinterface.
Step 15	**service-policy** {**input** | **output**} *parent-policy-name*	Applies the parent policy to the subinterface. **input** indicates to apply the service policy to inbound traffic. **output** indicates to apply the service policy to outbound traffic. *parent-policy-name* is the name of the parent policy map. **Note** When congestion occurs, the class queues receive bandwidth according to the specified class-level bandwidth-remaining ratios.

Configuration Examples for Distribution of Remaining Bandwidth Using Ratio

This section provides the following configuration examples:

Configuring Bandwidth-Remaining Ratios on Ethernet Subinterfaces: Example

The following example shows how to configure bandwidth-remaining ratios on an Ethernet subinterface using a hierarchical policy. In the example, Gigabit Ethernet subinterface 1/0/0.1 is shaped to 100 Mbps. During congestion, the router uses the bandwidth-remaining ratio of 10 to determine the amount of excess bandwidth (unused by priority traffic) to allocate to the non-priority traffic on subinterface 1/0/0.1, relative to the other subinterface-level and class-level queues on the 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 100000000
        service-policy Child
!
interface GigabitEthernet1/0/0.1
    encapsulation dot1Q 100
    ip address 10.1.0.1 255.255.255.0
    service-policy output Parent
```

Configuring Bandwidth-Remaining Ratios on ATM Subinterfaces: Example

The following example shows how to differentiate one ATM PVC from another during congestion by using bandwidth-remaining ratios. In the example, during periods of congestion in which the traffic on all PVCs on the interface exceeds the interface speed, the router uses the configured bandwidth-remaining ratio of 10 to determine the amount of excess (unused by priority traffic) bandwidth to allocate to non-priority traffic on PVC 0/200, relative to the other ATM PVCs configured on the interface.

```
policy-map Child
    class precedence_0
        bandwidth 100
    class precedence_1
        bandwidth 10000
!
policy-map Parent
    class class-default
        bandwidth remaining ratio 10
        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
```

Note If PVC 98/204 is configured on the same interface as PVC 0/200 and with a bandwidth-remaining ratio of 1, during times of congestion PVC 0/200 would have 10 times more bandwidth available to it for non-priority traffic than PVC 98/204 would have.

Configuring Bandwidth-Remaining Ratios on Class Queues: Example

In the following sample configuration, the vlan10_policy is applied on the subinterface Gigabit Ethernet 1/0/0.10 and the vlan20_policy is applied on the subinterface Gigabit Ethernet 1/0/0.20. During congestion on the interface, subinterface GE 1/0/0.20 has 10 times more available bandwidth than subinterface GE1/0/0.10 because the bandwidth-remaining ratio for subinterface GE 1/0/0.20 is 10 times more than the bandwidth-remaining ratio for subinterface 1/0/0.10: 100 on subinterface 1/0/0.20 and 10 on subinterface 1/0/0.10.

When congestion occurs within a subinterface level, the class queues receive bandwidth according to the class-level bandwidth-remaining ratios. In the example, the bandwidth for classes precedence_0, precedence_1, and precedence_2 is allocated based on the bandwidth-remaining ratios of the classes: 20, 40, and 60, respectively.

```
policy-map child-policy
    class precedence_0
        shape average 500000
        bandwidth remaining ratio 20 <---- Class-level ratio
    class precedence_1
        shape average 500000
        bandwidth remaining ratio 40 <---- Class-level ratio
    class precedence_2
        shape average 500000
        bandwidth remaining ratio 60 <---- Class-level ratio
!
policy-map vlan10_policy
    class class-default
        shape average 1000000
        bandwidth remaining ratio 10 <---- Subinterface-level ratio
        service-policy child-policy
!
policy-map vlan20_policy
    class class-default
        shape average 1000000
        bandwidth remaining ratio 100 <---- Subinterface-level ratio
        service-policy child_policy
!
!
interface GigabitEthernet 1/0/0.10
    encapsulation dot1q 10
    service-policy output vlan10_policy
!
interface GigabitEthernet 1/0/0.20
    encapsulation dot1q 20
    service-policy output vlan20_policy
```

Verifying Bandwidth Remaining Ratios: Example

The following sample output from the **show policy-map interface** command indicates that bandwidth-remaining ratios are configured on class-level queues in the policy maps named vlan10_policy and child_policy, which are attached to the Gigabit Ethernet subinterface 1/0/0.10.

```
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
```

The following sample output from the **show policy-map interface** command indicates that bandwidth-remaining ratios are configured on class-level queues in the policy maps named vlan20_policy and child_policy, which are attached to the Gigabit Ethernet subinterface 1/0/0.20.

```
Router# show policy-map interface GigabitEthernet1/0/0.20

  Service-policy output: vlan20_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 100

      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
```

The following sample output from the **show policy-map** command indicates that a bandwidth-remaining ratio of 10 is configured on the parent class-default class of the policy map named vlan10_policy.

```
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
```

The following sample output from the **show policy-map** command indicates that a bandwidth-remaining ratio of 100 is configured on the parent class-default class of the policy map named vlan20_policy. During congestion, the scheduler allocates the subinterface Gigabit Ethernet 1/0/0.20 10 times the bandwidth that it allocates subinterface Gigabit Ethernet 1/0/0.10.

```
Router# show policy-map vlan20_policy
  Policy Map vlan20_policy
    Class class-default
      Average Rate Traffic Shaping
      cir 1000000 (bps)
      bandwidth remaining ratio 100
      service-policy child_policy
```

The following sample output from the **show policy-map** command indicates that a bandwidth-remaining ratio of 20, 40, and 60 is configured on the class queues precedence_0, precedence_1, and precedence_2, respectively.

```
Router# show policy-map child_policy
  Policy Map child_policy
    Class precedence_0
      Average Rate Traffic Shaping
      cir 500000 (bps)
      bandwidth remaining ratio 20
    Class precedence_1
      Average Rate Traffic Shaping
      cir 500000 (bps)
      bandwidth remaining ratio 40
    Class precedence_2
      Average Rate Traffic Shaping
      cir 500000 (bps)
      bandwidth remaining ratio 60
```

Related Documentation

This section provides hyperlinks to additional Cisco documentation for the features discussed in this chapter. To display the documentation, click the document title or a section of the document highlighted in blue. When appropriate, paths to applicable sections are listed below the documentation title.

Feature	Related Documentation
bandwidth command	*Cisco IOS Quality of Service Solutions Command Reference, Release 12.3 T* Quality of Service Commands: A through F > bandwidth (policy-map class) Comparing the bandwidth and priority Commands of a QoS Service Policy
bandwidth remaining ratio	Distribution of Remaining Bandwidth Using Ratio, Release 12.2(31)SB2 feature module
Class maps	*Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2* Part 8: Modular Quality of Service Command-Line Interface > Configuring the Modular Quality of Service Command-Line Interface > Modular QoS CLI Configuration Task List > Creating a Traffic Class Cisco IOS Quality of Service Solutions Command Reference, Release 12.2 access-list rate-limit -- fair-queue (WFQ) > class-map command
Hierarchical queuing framework	Cisco 10000 Series Router Quality of Service Configuration Guide Hierarchical Scheduling and Queuing
Policing	Comparing Traffic Shaping and Traffic Policing for Bandwidth Limiting
Policy maps	*Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2* Part 8: Modular Quality of Service Command-Line Interface > Configuring the Modular Quality of Service Command-Line Interface > Modular QoS CLI Configuration Task List > Creating a Traffic Policy *Cisco IOS Quality of Service Solutions Command Reference, Release 12.2* policy map - qos preclassify > policy-map command
QoS service policies	*QoS Configuration and Monitoring, Creating Time-of-Day QoS Service Policies* tech note *QoS Configuration and Monitoring, Monitoring Voice over IP Quality of Service* tech note *Site-to-Site MPLS VPN Solution for Service Providers, Service Provider Quality-of-Service Overview* tech note
Shaping	Comparing Traffic Shaping and Traffic Policing for Bandwidth Limiting
Three-level scheduler	Cisco 10000 Series Router Quality of Service Configuration Guide Hierarchical Scheduling and Queuing