



QoS: IP to ATM Class of Service Configuration Guide Cisco IOS Release 12.2SX

Americas Headquarters

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



CONTENTS

IP to ATM Class of Service Overview	1
Finding Feature Information	1
About IP to ATM CoS	1
Single ATM VC Support	2
VC Bundle Support and Bundle Management	2
Per-VC LLQ WFQ and CBWFQ Support	4
Why Use IP to ATM CoS	5
Benefits	5
IP to ATM CoS Features	5
Congestion Avoidance	6
Bumping and ATM VC Bundles	7
Restrictions	8
Configuring IP to ATM Class of Service	9
Finding Feature Information	9
IP to ATM CoS on a Single ATM VC Configuration Task List	9
Defining the WRED Parameter Group	10
Configuring the WRED Parameter Group	10
Displaying the WRED Parameters	10
Displaying the Queueing Statistics	10
IP to ATM CoS on an ATM Bundle Configuration Task List	11
Creating a VC Bundle	11
Applying Bundle-Level Parameters	11
Configuring Bundle-Level Parameters	11
Configuring VC Class Parameters to Apply to a Bundle	12
Attaching a Class to a Bundle	12
Committing a VC to a Bundle	13
Applying Parameters to Individual VCs	13
Configuring a VC Bundle Member Directly	13
Configuring VC Class Parameters to Apply to a VC Bundle Member	13
Applying a VC Class to a Discrete VC Bundle Member	14
Configuring a VC Not to Accept Bumped Traffic	14

Monitoring and Maintaining VC Bundles and Their VC Members	14
Per-VC WFQ and CBWFQ Configuration Task List	15
Configuring Class-Based Weighted Fair Queueing	15
Attaching a Service Policy and Enabling CBWFQ for a VC	16
Attaching a Policy Map to a Standalone VC and Enabling CBWFQ	16
Attaching a Policy Map to an Individual VC and Enabling CBWFQ	16
Configuring a VC to Use Flow-Based WFQ	16
Attaching a Policy Map to a Standalone VC and Enabling WFQ	17
Attaching a Policy Map to an Individual VC and Enabling WFQ	18
Monitoring per-VC WFQ and CBWFQ	18
Enabling Logging of Error Messages to the Console	18
IP to ATM CoS Configuration Examples	18
Example Single ATM VC with WRED Group and IP Precedence	18
Example VC Bundle Configuration Using a VC Class	19
Bundle-Class Class	19
Control-Class Class	19
Premium-Class Class	20
Priority-Class Class	20
Basic-Class Class	20
new-york Bundle	20
san-francisco Bundle	21
los-angeles Bundle	21
Example Per-VC WFQ and CBWFQ on a Standalone VC	22
Example Per-VC WFQ and CBWFQ on Bundle-Member VCs	23



IP to ATM Class of Service Overview

This module provides a high-level overview of IP to ATM Class of Service (CoS), a feature suite that maps QoS characteristics between IP and ATM.

- [Finding Feature Information, page 1](#)
- [About IP to ATM CoS, page 1](#)
- [Why Use IP to ATM CoS, page 5](#)
- [IP to ATM CoS Features, page 5](#)
- [Congestion Avoidance, page 6](#)
- [Bumping and ATM VC Bundles, page 7](#)
- [Restrictions, page 8](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

About IP to ATM CoS

The IP to ATM CoS feature implements a solution for coarse-grained mapping of QoS characteristics between IP and ATM, using Cisco Enhanced ATM port adapters (PA-A3) on Cisco 7200 and Cisco 7500 series routers. (This category of coarse-grained QoS is often referred to as CoS). The resulting feature makes it possible to support differential services in network service provider environments.

IP to ATM CoS is designed to provide a true working solution to class-based services, without the investment of new ATM network infrastructures. Now networks can offer different service classes (sometimes termed *differential service classes*) across the entire WAN, not just the routed portion. Mission-critical applications can be given exceptional service during periods of high network usage and congestion. In addition, noncritical traffic can be restricted in its network usage, which ensures greater QoS for more important traffic and user types.

The IP to ATM CoS feature is supported on Cisco 2600, Cisco 3600, Cisco 7200, and Cisco 7500 series routers equipped with the following hardware:

- Cisco 2600 and Cisco 3600 series: ATM OC-3, T1 IMA, or E1 IMA port adapter
- Cisco 7200 series:
 - NPE-200 or higher (NPE-300 recommended for per-virtual circuit (VC) class-based weighted fair queueing (CBWFQ))
 - One of the following Enhanced ATM port adapters (PA-A3): T3, E3, DS3, or OC-3
- Cisco 7500 series:
 - VIP2-50
 - One of the following Enhanced ATM port adapters (PA-A3): T3, E3, DS3, or OC-3

IP to ATM CoS supports configuration of the following features:

- [Single ATM VC Support, page 2](#)
- [VC Bundle Support and Bundle Management, page 2](#)
- [Per-VC LLQ WFQ and CBWFQ Support, page 4](#)


Single ATM VC Support

IP to ATM CoS support for a single ATM VC allows network managers to use existing features, such as committed access rate (CAR) or policy-based routing (PBR), to classify and mark different IP traffic by modifying the IP Precedence field in the IP version 4 (IPv4) packet header. Subsequently, Weighted Random Early Detection (WRED) or distributed WRED (DWRED) can be configured on a per-VC basis so that the IP traffic is subject to different drop probabilities (and therefore priorities) as IP traffic coming into a router competes for bandwidth on a particular VC.

Enhanced ATM port adapters (PA-A3) provide the ability to shape traffic on each VC according to the ATM service category and traffic parameters employed. When you use the IP to ATM CoS feature, congestion is managed entirely at the IP layer by WRED running on the routers at the edge of the ATM network.

The figure below illustrates the IP to ATM CoS support for a single ATM VC.


Figure 1: Single ATM Circuit Class



VC Bundle Support and Bundle Management

ATM VC bundle management allows you to configure multiple VCs that have different QoS characteristics between any pair of ATM-connected routers. As shown in the figure below, these VCs are grouped in a bundle and are referred to as bundle members.

Figure 2: ATM VC Bundle




ATM VC bundle management allows you to define an ATM VC bundle and add VCs to it. Each VC of a bundle has its own ATM traffic class and ATM traffic parameters. You can apply attributes and characteristics to discrete VC bundle members or you can apply them collectively at the bundle level.

Using VC bundles, you can create differentiated service by flexibly distributing IP precedence levels over the different VC bundle members. You can map a single precedence level or a range of levels to each discrete VC in the bundle, thereby enabling individual VCs in the bundle to carry packets marked with different precedence levels. You can use WRED (or DWRED) to further differentiate service across traffic that has different IP precedences but that uses the same VC in a bundle.

To determine which VC in the bundle to use to forward a packet to its destination, the ATM VC bundle management software matches precedence levels between packets and VCs (see the figure below). IP traffic is sent to the next hop address for the bundle because all VCs in a bundle share the same destination, but the VC used to carry a packet depends on the value set for that packet in the IP Precedence bits of the type of service (ToS) byte of its header. The ATM VC bundle management software matches the IP precedence of the packet to the IP Precedence value or range of values assigned to a VC, sending the packet out on the appropriate VC. Moreover, the ATM VC bundle management feature allows you to configure how traffic will be redirected when the VC the packet was matched to goes down. The figure below illustrates how the ATM VC bundle management software determines which permanent virtual

circuit (PVC) bundle member to use to carry a packet and how WRED (or DWRED) is used to differentiate traffic on the same VC.

Figure 3: ATM VC Bundle PVC Selection for Packet Transfer



The support of multiple parallel ATM VCs allows you to create stronger service differentiation at the IP layer. For instance, you might want to provide IP traffic belonging to real-time CoS (such as Voice over IP traffic) on an ATM VC with strict constraints (constant bit rate (CBR) or variable bit rate real-time (VBR-rt), for example), while transporting traffic other than real-time traffic over a more elastic ATM available bit rate (ABR) PVC. Using a configuration such as this would allow you to fully utilize your network capacity. You could also elect to transport best-effort IP traffic over an unspecified bit rate (UBR) PVC--UBR is effectively the ATM version of best-effort service.

Per-VC LLQ WFQ and CBWFQ Support

The IP to ATM CoS feature allows you to apply a policy map to a VC to specify a service policy for that VC so that all traffic sent on that VC is categorized according to the classes and their match criteria defined by the service policy. In other words, IP to ATM CoS takes the functionality defined for standard LLQ, WFQ, and CBWFQ and makes it available for application and use at the discrete VC level.

IP to ATM CoS allows you to configure a single, standalone VC or individual VCs belonging to a bundle. You also can configure collectively all VCs belonging to a bundle. However, for per-VC LLQ, WFQ and CBWFQ, you can configure individual VCs only. That is, you can configure a standalone VC or a VC that belongs to a bundle, but you cannot use per-VC LLQ, WFQ and CBWFQ to configure a bundle of VCs collectively.

Per-VC LLQ, WFQ and CBWFQ allows you to differentiate the use of individual VCs within a bundle. For instance, you can apply one service policy to one VC belonging to a VC bundle and apply a different service policy to another VC belonging to the same bundle. You can also apply the same policy map to multiple VCs--whether standalone or bundle members--but each VC can have only one service policy. To concatenate service policies, you must create a third policy map and include in it all the classes that you want to use from policy maps you would have concatenated.

The following is a summary of how you configure a VC to use CBWFQ:

- You define traffic classes to specify the classification policy (class maps). This process determines how many types of packets are to be differentiated from one another.
- You configure policy maps containing classes that specify the policy for each traffic class.
- You attach a policy map to a VC that uses IP to ATM CoS to specify the service policy for the VC.

To apply flow-based WFQ on a per-VC basis, you configure WFQ in the predefined CBWFQ default class, which is called class-default, but you do not ascribe bandwidth to the default class.

Why Use IP to ATM CoS

Internet service classes can be identified and sorted within the router network. But as traffic traverses the wide-area ATM fabric, the relative ATM class definitions are not equivalent, and a traffic type may be treated differently in the ATM switching fabric than in the router network; mission-critical applications or data could be dropped during times of network congestion.

The IP to ATM CoS feature uses the Cisco Enhanced ATM port adapter (PA-A3) on Cisco 7500 and Cisco 7200 series routers to provide the ability to map IP CoS and ATM QoS, extending the capability previously available only for IP networks; differentiated services are preserved through the ATM network.

- [Benefits, page 5](#)

Benefits

Here are some benefits of using IP to ATM CoS:

- Ensures effective differential classes over IP and traditional ATM networks. For instance, the VC bundle management feature provides for differentiated QoS by allowing for the coexistence of multiple VCs with different QoS characteristics from the same source to the same destination.
- Uses existing ATM infrastructures.
- Implements solutions for coarse-grained mapping of QoS characteristics called CoS between IP and ATM.
- Employs a high-performance design benefiting from distributed processing on the Cisco 7500 series routers and Versatile Interface Processor (VIP).
- Uses the Cisco Enhanced ATM port adapter (PA-A3), which supports traffic shaping and has rich ATM Service Category support. This port adapter (PA) is supported on the Cisco 7500+VIP and Cisco 7200 series routers.
- Provides per-VC queueing on the PA, per-VC back pressure, and per-VC WRED VIP queueing, which bring stability to a network by ensuring that system packets such as Border Gateway Protocol (BGP) and Intermediate System-to-Intermediate System (IS-IS) are never dropped.
- Provides flexible management of the VC bundle on PVC failure.
- Provides CBWFQ functionality at the VC level.

IP to ATM CoS Features

IP to ATM CoS includes the following features:

- Per-VC queueing infrastructure. This feature enables queues to be maintained on a per-VC basis. Packets are queued and dequeued based on the back pressure from the PA. Use of a queue per VC prevents one or more congested VCs from affecting the traffic flow on other VCs that are not congested.

- Per-VC WRED (or DWRED). This feature applies the WRED algorithm independently to each per-VC queue. The WRED parameters are configurable on a per-VC basis so that congestion management can be configured as appropriate for each VC.
- Per-VC WRED (or DWRED) statistics. This feature maintains per-flow and per-VC statistics based on IP precedence.
- Per-VC LLQ, WFQ and CBWFQ. This feature allows you to apply CBWFQ functionality--normally applicable at the interface or subinterface levels only--to an individual VC configured for IP to ATM CoS. You can use this feature to apply either CBWFQ or flow-based WFQ on a per-VC basis.
- Per-VC traffic policing. This feature allows you to police traffic within a traffic policy, per-VC.

Congestion Avoidance

For each VC that is created on the Enhanced ATM port adapter (PA-A3), the PA allocates some of the buffers from its buffer pool to that VC in order to create a queue for that VC.


The use of per-VC queues ensures that a direct relationship exists between the outgoing ATM VC and the IP packets to be forwarded on that queue. This mechanism establishes a packet queue for each outgoing ATM VC. In this manner, should an ATM VC become congested, only the packet queue associated with that VC will begin to fill. If the queue overfills, then all other queues remain unaffected. Such a mechanism ensures that an individual VC cannot consume all of the resources of the router should only one of its outgoing VCs be congested or underprovisioned.

Queues for buffering more packets for a particular VC are created in the Layer 3 processor system and are mapped one-to-one to the per-VC queues on the PA. When the PA per-VC queues become congested, they signal back pressure to the Layer 3 processor; the Layer 3 processor can then continue to buffer packets for that VC in the corresponding Layer 3 queue. Furthermore, because the Layer 3 queues are accessible by the Layer 3 processor, a user can run flexible software scheduling algorithms on those queues.

When you transport data over ATM fabrics, it is essential that decisions to discard data (because of insufficient network resources or congestion) be made at the packet level. To do otherwise would be to send incomplete data packets into the ATM fabric, causing the packets to be discarded by either the ATM switched fabric (if it is equipped with early packet discard) or at the remote end where the packet will be reassembled and found to be incomplete.

To initiate effective congestion management techniques, IP to ATM CoS uses per-VC WRED (or DWRED). Per-VC WRED (or DWRED) selectively places TCP sessions in slow start mode to ensure higher aggregate throughput under congestion. The figure below shows low priority packets being dropped on VC1 because VC1 is congested. In this example, VC2 is not congested and all packets, regardless of priority, are sent.

Figure 4: Traffic Congestion with IP to ATM CoS and Per-VC WRED



Running the WRED algorithm independently on each per-VC queue provides differentiated QoS to traffic of different IP Precedence values.


Bumping and ATM VC Bundles

The ATM VC bundle is designed to behave as a single routing link to the destination router while managing the integrity of its group of circuits. The integrity of each circuit is maintained through individual monitoring. Should a circuit fail, appropriate action is taken, in the form of circuit bumping or bundle disabling.

VC integrity is maintained through ATM Operation, Administration, and Maintenance (OAM) polling mechanisms. These mechanisms will determine whether a VC is unavailable or severely congested. Should an individual circuit become unavailable, then the device consults a preset series of rules to determine the course of action to take next. These rules are defined by the Internet service provider (ISP) through configuration parameters.

The figure below conceptualizes a failed VC bundle member whose failure calls into effect the configured bumping rules.

Figure 5: VC Bundle Member Circuit Failure Enacting Bumping Rules



In the event of failure, the router responds with one of two methods. The first method dynamically assigns the traffic bound on the failed VC to an alternative VC, which is termed *circuit bumping*. Bumped traffic is then shared on an existing in-service VC. Traffic typically would be bumped from a higher class to a lower one, although it need not be. For example, should the premium, or first class, data circuit become unavailable, then all premium users would share the second class or general circuit. Preference would then be given to the premium traffic within this shared circuit.

The second method is to declare all circuits of the bundle to be down. In effect, the device is declaring the routed bundle inactive and asking the routing layer to search for an alternate.

The determination of whether to bump or whether to declare the bundle inactive is predefined by the network provider when administering the network configuration.

Restrictions

The following restrictions apply for IP to ATM CoS:

- IP to ATM CoS supports only PVCs:
 - For PVC connections, it supports multipoint and point-to-point subinterfaces.
 - For PVC encapsulations, it supports only ATM adaptation layer (AAL5), Subnetwork Access Protocol (SNAP), and multiplex device (mux) interfaces.
- IP to ATM CoS does not allow point-to-multipoint VCs in the bundle. All VCs share the same source and destination (target) addresses.
- IP to ATM CoS does not work with the ATM Interface Processor (AIP) and the ATM port adapter (PA-A1).

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned 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. (1005R)

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



Configuring IP to ATM Class of Service

This module describes the tasks for configuring the IP to ATM Class of Service (CoS), a feature suite that maps QoS characteristics between IP and ATM.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS and Catalyst OS software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

- [Finding Feature Information, page 9](#)
- [IP to ATM CoS on a Single ATM VC Configuration Task List, page 9](#)
- [IP to ATM CoS on an ATM Bundle Configuration Task List, page 11](#)
- [Per-VC WFQ and CBWFQ Configuration Task List, page 15](#)
- [IP to ATM CoS Configuration Examples, page 18](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

IP to ATM CoS on a Single ATM VC Configuration Task List

To configure IP to ATM CoS for a single ATM virtual circuit (VC), perform the tasks described in the following sections. The tasks in the first two sections are required; the tasks in the remaining sections are optional.

The IP to ATM CoS feature requires ATM permanent virtual circuit (PVC) management.

- [Defining the WRED Parameter Group, page 10](#)
- [Configuring the WRED Parameter Group, page 10](#)
- [Displaying the WRED Parameters, page 10](#)
- [Displaying the Queueing Statistics, page 10](#)

Defining the WRED Parameter Group

Command	Purpose
Router(config)# random-detect-group <i>group-name</i>	Defines the WRED or VIP-distributed WRED (DWRED) parameter group.

Configuring the WRED Parameter Group

SUMMARY STEPS

1. Router(config)# **random-detect-group** *group-name*
2. Router(config)# **exponential-weighting-constant** *exponent* Router(config)# **precedence** *precedence min-threshold max-threshold mark-probability-denominator*

DETAILED STEPS

	Command or Action	Purpose
Step 1	Router(config)# random-detect-group <i>group-name</i>	Specifies the WRED or DWRED parameter group.
Step 2	Router(config)# exponential-weighting-constant <i>exponent</i> Router(config)# precedence <i>precedence min-threshold max-threshold mark-probability-denominator</i>	Configures the exponential weight factor for the average queue size calculation for the specified WRED or DWRED parameter group. or Configures the specified WRED or DWRED parameter group for a particular IP precedence.

Displaying the WRED Parameters

Command	Purpose
Router# show queueing random-detect [interface <i>atm_subinterface</i> [vc [[<i>vpi/</i>] <i>vci</i>]]]	Displays the parameters of every VC with WRED or DWRED enabled on the specified ATM subinterface.

Displaying the Queueing Statistics

Command	Purpose
Router# show queueing interface <i>interface-number</i> [vc [[<i>vpi/</i>] <i>vci</i>]]	Displays the queueing statistics of a specific VC on an interface.

IP to ATM CoS on an ATM Bundle Configuration Task List

To configure IP to ATM CoS on an ATM bundle, perform the tasks in the following sections.

The IP to ATM CoS feature requires ATM PVC management.

- [Creating a VC Bundle, page 11](#)
- [Applying Bundle-Level Parameters, page 11](#)
- [Committing a VC to a Bundle, page 13](#)
- [Applying Parameters to Individual VCs, page 13](#)
- [Configuring a VC Not to Accept Bumped Traffic, page 14](#)
- [Monitoring and Maintaining VC Bundles and Their VC Members, page 14](#)

Creating a VC Bundle

Command	Purpose
Router(config-subif)# bundle <i>bundle-name</i>	Creates the specified bundle and enters bundle configuration mode.

Applying Bundle-Level Parameters

- [Configuring Bundle-Level Parameters, page 11](#)
- [Configuring VC Class Parameters to Apply to a Bundle, page 12](#)
- [Attaching a Class to a Bundle, page 12](#)

Configuring Bundle-Level Parameters

Command	Purpose
Router(config-atm-bundle)# protocol <i>protocol</i> { <i>protocol-address</i> inarp } [[no] broadcast]	Configures a static map or enables Inverse Address Resolution Protocol (Inverse ARP) or Inverse ARP broadcasts for the bundle. Note Bundle-level parameters can be applied either by assigning VC classes or by directly applying them to the bundle. Parameters applied through a VC class assigned to the bundle are superseded by those applied at the bundle level. Bundle-level parameters are superseded by parameters applied to an individual VC.
Router(config-atm-bundle)# encapsulation <i>aal-encap</i>	Configures the ATM adaptation layer (AAL) and encapsulation type for the bundle.
Router(config-atm-bundle)# inarp <i>minutes</i>	Configures the Inverse ARP time period for all VC bundle members.

Command	Purpose
Router(config-atm-bundle)# broadcast	Enables broadcast forwarding for all VC bundle members.
Router(config-atm-bundle)# oam retry up-count down-count retry frequency	Configures the VC bundle parameters related to operation, administration, and maintenance (OAM) management.
Router(config-atm-bundle)# oam-bundle [manage] [frequency]	Enables end-to-end F5 OAM loopback cell generation and OAM management for all VCs in the bundle.

Configuring VC Class Parameters to Apply to a Bundle

Command	Purpose
Router(config-vc-class)# oam-bundle [manage] [frequency]	Enables end-to-end F5 OAM loopback cell generation and OAM management for all VCs in the bundle. Note Use of a VC class allows you to configure a bundle applying multiple attributes to it at once because you apply the class itself to the bundle. Use of a class allows you to generalize a parameter across all VCs, after which (for some parameters) you can modify that parameter for individual VCs. (See the section "Applying Parameters to Individual VCs" for more information.)

Attaching a Class to a Bundle

Command	Purpose
Router(config-atm-bundle)# class-bundle vc-class-name	Configures a bundle with the bundle-level commands contained in the specified VC class. Note Parameters set through bundle-level commands contained in the VC class are applied to the bundle and all of its VC members. Bundle-level parameters applied through commands configured directly on the bundle supersede those applied through a VC class. Some bundle-level parameters applied through a VC class or directly to the bundle can be superseded by commands that you directly apply to individual VCs in bundle-vc configuration mode.

Committing a VC to a Bundle

Command	Purpose
Router(config-atm-bundle)# pvc-bundle <i>pvc-name</i> [<i>vpi/</i>] [<i>vci</i>]	Adds the specified VC to the bundle and enters bundle-vc configuration mode in order to configure the specified VC bundle member.

Applying Parameters to Individual VCs

- [Configuring a VC Bundle Member Directly, page 13](#)
- [Configuring VC Class Parameters to Apply to a VC Bundle Member, page 13](#)
- [Applying a VC Class to a Discrete VC Bundle Member, page 14](#)

Configuring a VC Bundle Member Directly

Command	Purpose
Router(config-if-atm-member)# ubr <i>output-pcr</i> [<i>input-pcr</i>]	Configures the VC for unspecified bit rate (UBR) QoS and specifies the output peak cell rate (PCR) for it.
Router(config-if-atm-member)# ubr+ <i>output-pcr output-mcr</i> [<i>input-pcr</i>] [<i>input-mcr</i>]	Configures the VC for UBR QoS and specifies the output PCR and output minimum guaranteed cell rate for it.
Router(config-if-atm-member)# vbr-nrt <i>output-pcr output-scr output-mbs</i> [<i>input-pcr</i>] [<i>input-scr</i>] [<i>input-mbs</i>]	Configures the VC for variable bit rate nonreal-time (VBR-nrt) QoS and specifies the output PCR, output sustainable cell rate, and output maximum burst cell size for it.
Router(config-if-atm-member)# precedence [other <i>range</i>]	Configures the precedence levels for the VC.
Router(config-if-atm-member)# bump { implicit explicit <i>precedence-level</i> traffic }	Configures the bumping rules for the VC.
Router(config-if-atm-member)# protect { group vc }	Configures the VC to belong to the protected group of the bundle or to be an individually protected VC bundle member.

Configuring VC Class Parameters to Apply to a VC Bundle Member

Command	Purpose
Router(config-vc-class)# bump { implicit explicit <i>precedence-level</i> traffic }	Specifies the bumping rules for the VC member to which the class is applied. These rules determine to which VC in the bundle traffic is directed when the carrier VC bundle member goes down. Note You can also add the following commands to a VC class to be used to configure a VC bundle member: ubr , ubr+ , and vbr-nrt . When a VC is a member of a VC bundle, the following commands cannot be used in vc-class mode to configure the VC: encapsulation , protocol , inarp , and broadcast . These commands are useful only at the bundle level, not the bundle member level. Configuration for an individual VC overrides the collective configuration applied to all VC bundle members through application of a VC class to the bundle.
Router(config-vc-class)# precedence <i>precedence min-threshold max-threshold mark-probability-denominator</i>	Defines precedence levels for the VC member to which the class is applied.
Router(config-vc-class)# protect { group vc }	Configures the VC as a member of the protected group of the bundle or as an individually protected VC.

Applying a VC Class to a Discrete VC Bundle Member

Command	Purpose
Router(config-if-atm-member)# class-vc <i>vc-class -name</i>	Assigns a VC class to a VC bundle member.

Configuring a VC Not to Accept Bumped Traffic

Command	Purpose
Router(config-if-atm-member)# no bump traffic	Configures the VC not to accept any bumped traffic that would otherwise be redirected to it.

Monitoring and Maintaining VC Bundles and Their VC Members

Command	Purpose
Router# show atm bundle <i>bundle-name</i>	Displays the bundle attributes assigned to each bundle VC member and the current working status of the VC members.
Router# show atm bundle statistics [detail]	Displays statistics or detailed statistics on the specified bundle.
Router# show atm map	Displays a list of all configured ATM static maps to remote hosts on an ATM network and on ATM bundle maps.
Router# debug atm bundle errors	Displays information on bundle errors.
Router# debug atm bundle events	Displays a record of bundle events.

Per-VC WFQ and CBWFQ Configuration Task List

To configure IP to ATM CoS for per-VC WFQ and CBWFQ, perform the tasks described in the following sections.

The IP to ATM CoS feature requires ATM PVC management.

- [Configuring Class-Based Weighted Fair Queueing, page 15](#)
- [Attaching a Service Policy and Enabling CBWFQ for a VC, page 16](#)
- [Configuring a VC to Use Flow-Based WFQ, page 16](#)
- [Monitoring per-VC WFQ and CBWFQ, page 18](#)
- [Enabling Logging of Error Messages to the Console, page 18](#)

Configuring Class-Based Weighted Fair Queueing

Before configuring CBWFQ for a VC, you must perform the following tasks using standard CBWFQ commands:

- Create one or more classes to be used to classify traffic sent across the VC
- Define a policy map containing the classes to be used as the service policy



Note

You can configure class policies for as many classes as are defined on the router, up to the maximum of 64. However, the total amount of bandwidth allocated for all classes included in a policy map to be attached to a VC must not exceed 75 percent of the available bandwidth of the VC. The remaining 25 percent of available bandwidth is used for encapsulation, such as the ATM cell overhead (also referred to as ATM cell tax), routing and best-effort traffic, and other functions that assume overhead. For more information on bandwidth allocation, see the "Congestion Management Overview" module.

Because CBWFQ gives you minimum bandwidth guarantee, you can only apply CBWFQ to VCs having these classes of service: available bit rate (ABR) and variable bit rate (VBR). You cannot apply per-VC WFQ and CBWFQ to UBR and unspecified bit rate plus (UBR+) VCs because both of these service classes

are best-effort classes that do not guarantee minimum bandwidth. When CBWFQ is enabled for a VC, all classes configured as part of the service policy are installed in the fair queueing system.

In addition to configuring CBWFQ at the VC level, the IP to ATM CoS feature allows you to configure flow-based WFQ at the VC level. Because flow-based WFQ gives you best-effort class of service--that is, it does not guarantee minimum bandwidth--you can configure per-VC WFQ for all types of CoS VCs: ABR, VBR, UBR, and UBR+.

Per-VC WFQ uses the class-default class. Therefore, to configure per-VC WFQ, you must first create a policy map and configure the class-default class. (You need not create the class-default class, which is predefined, but you must configure it.) For per-VC WFQ, the class-default class must be configured with the **fair-queue** policy-map class configuration command.

In addition to configuring the **fair-queue** policy-map class configuration command, you can configure the default class with either the **queue-limit** command or the **random-detect** command, but not both. Moreover, if you want the default class to use flow-based WFQ, you cannot configure the default class with the **bandwidth** policy-map class configuration command--to do so would disqualify the default class as flow-based WFQ, and therefore limit application of the service policy containing the class to ABR and VBR VCs.

Attaching a Service Policy and Enabling CBWFQ for a VC

- [Attaching a Policy Map to a Standalone VC and Enabling CBWFQ, page 16](#)
- [Attaching a Policy Map to an Individual VC and Enabling CBWFQ, page 16](#)

Attaching a Policy Map to a Standalone VC and Enabling CBWFQ

Command	Purpose
Router(config-if-atm-vc)# service-policy output <i>policy-map</i>	Enables CBWFQ and attaches the specified service policy map to the VC being created or modified.

Attaching a Policy Map to an Individual VC and Enabling CBWFQ

Command	Purpose
Router(config-if-atm-member)# service-policy output <i>policy-map</i>	Enables CBWFQ and attaches the specified service policy map to the VC being created or modified.



Note

The **service-policy output** and **random-detect-group** commands are mutually exclusive; you cannot apply a WRED group to a VC for which you have enabled CBWFQ through application of a service policy. Moreover, before you can configure one command, you must disable the other if it is configured.

Configuring a VC to Use Flow-Based WFQ

SUMMARY STEPS

1. Router(config)# **policy-map** *policy-map*
2. Router(config-pmap)# **class class-default** *default-class-name*
3. Router(config-pmap-c)# **fair-queue** *number-of-dynamic-queues*
4. Do one of the following:
 - Router(config-pmap-c)# **queue-limit** *number-of-packets*
 -
 -
 - Router(config-pmap-c)# **random-detect**

DETAILED STEPS

	Command or Action	Purpose
Step 1	Router(config)# policy-map <i>policy-map</i>	Specifies the name of the policy map to be created or modified.
Step 2	Router(config-pmap)# class class-default <i>default-class-name</i>	Specifies the default class so that you can configure or modify its policy. Note You can include other classes in the same policy map as the one that contains the flow-based WFQ class. Packets not otherwise matched are selected by the default class-default class match criteria.
Step 3	Router(config-pmap-c)# fair-queue <i>number-of-dynamic-queues</i>	Specifies the number of dynamic queues to be reserved for use by flow-based WFQ running on the default class. Note By default--that is, even if you do not configure the class-default class with the fair-queue policy-map class configuration command and you do not configure it with the bandwidth policy-map class configuration command--the default class is defined as flow-based WFQ.
Step 4	Do one of the following: <ul style="list-style-type: none"> • Router(config-pmap-c)# queue-limit <i>number-of-packets</i> • • Router(config-pmap-c)# random-detect 	Specifies the maximum number of packets that can be queued for the class. Enables WRED. The class policy will drop packets using WRED instead of tail drop.

- [Attaching a Policy Map to a Standalone VC and Enabling WFQ, page 17](#)
- [Attaching a Policy Map to an Individual VC and Enabling WFQ, page 18](#)

Attaching a Policy Map to a Standalone VC and Enabling WFQ

Command	Purpose
Router(config-if-atm-vc)# service-policy output <i>policy-map</i>	Enables WFQ for the VC by attaching the specified policy map containing the class-default class to the VC being created or modified.

Attaching a Policy Map to an Individual VC and Enabling WFQ

Command	Purpose
Router(config-if-atm-member)# service-policy output <i>policy-map</i>	Enables WFQ for the VC bundle member by attaching the specified policy map containing the class-default class to the VC bundle member.

Monitoring per-VC WFQ and CBWFQ

Command	Purpose
Router# show queue <i>interface-name interface-number</i> [vc [<i>vpi/</i>] <i>vci</i>]	Displays the contents of packets inside a queue for a particular interface or VC.
Router# show queueing interface <i>interface-number</i> [vc [<i>vpi/</i>] <i>vci</i>]	Displays the queueing statistics of a specific VC on an interface.

Enabling Logging of Error Messages to the Console

Command	Purpose
Router(config)# logging console <i>level</i>	Limits messages logged to the console based on severity.

IP to ATM CoS Configuration Examples

- [Example Single ATM VC with WRED Group and IP Precedence, page 18](#)
- [Example VC Bundle Configuration Using a VC Class, page 19](#)
- [Example Per-VC WFQ and CBWFQ on a Standalone VC, page 22](#)
- [Example Per-VC WFQ and CBWFQ on Bundle-Member VCs, page 23](#)

Example Single ATM VC with WRED Group and IP Precedence

The following example creates a PVC on an ATM interface and applies the WRED parameter group called sanjose to that PVC. Next, the IP Precedence values are configured for the WRED parameter group sanjose.

```
interface ATM1/1/0.46 multipoint
 ip address 200.126.186.2 255.255.255.0
 no ip mroute-cache
 shutdown
 pvc cisco 46
 encapsulation aal5nlpid
 random-detect attach sanjose
 !
 random-detect-group sanjose
```

```

precedence 0 200 1000 10
precedence 1 300 1000 10
precedence 2 400 1000 10
precedence 3 500 1000 10
precedence 4 600 1000 10
precedence 5 700 1000 10
precedence 6 800 1000 10
precedence 7 900 1000 10

```

Example VC Bundle Configuration Using a VC Class

This example configures VC bundle management on a router that uses Intermediate System-to-Intermediate System (IS-IS) as its IP routing protocol.

- [Bundle-Class Class, page 19](#)
- [Control-Class Class, page 19](#)
- [Premium-Class Class, page 20](#)
- [Priority-Class Class, page 20](#)
- [Basic-Class Class, page 20](#)
- [new-york Bundle, page 20](#)
- [san-francisco Bundle, page 21](#)
- [los-angeles Bundle, page 21](#)

Bundle-Class Class

At the outset, this configuration defines a VC class called `bundle-class` that includes commands that set VC parameters. When the class `bundle-class` is applied at the bundle level, these parameters are applied to all VCs that belong to the bundle. Note that any commands applied directly to an individual VC of a bundle in `bundle-vc` mode take precedence over commands applied globally at the bundle level. Taking into account hierarchy precedence rules, VCs belonging to any bundle to which the class `bundle-class` is applied will be characterized by these parameters: `aal5snap` encapsulation, broadcast on, use of Inverse Address Resolution Protocol (ARP) to resolve IP addresses, and operation, administration, and maintenance (OAM) enabled.

```

router isis
 net 49.0000.0000.0000.1111.00
vc-class atm bundle-class
 encapsulation aal5snap
 broadcast
 protocol ip inarp
 oam-bundle manage 3
 oam retry 4 3 10

```

The following sections of the configuration define VC classes that contain commands specifying parameters that can be applied to individual VCs in a bundle by assigning the class to that VC.

Control-Class Class

When the class called `control-class` is applied to a VC, the VC carries traffic whose IP Precedence level is 7. When the VC to which this class is assigned goes down, it takes the bundle down with it because this class makes the VC a protected one. The QoS type of a VC using this class is `vbr-nrt`.

```

vc-class atm control-class
 precedence 7
 protect vc
 vbr-nrt 10000 5000 32

```

Premium-Class Class

When the class called premium-class is applied to a VC, the VC carries traffic whose IP Precedence levels are 6 and 5. The VC does not allow other traffic to be bumped onto it. When the VC to which this class is applied goes down, its bumped traffic will be redirected to a VC whose IP Precedence level is 7. This class makes a VC a member of the protected group of the bundle. When all members of a protected group go down, the bundle goes down. The QoS type of a VC using this class is vbr-nrt.

```
vc-class atm premium-class
precedence 6-5
no bump traffic
protect group
bump explicitly 7
vbr-nrt 20000 10000 32
```

Priority-Class Class

When the class called priority-class is applied to a VC, the VC is configured to carry traffic with IP Precedence in the 4-2 range. The VC uses the implicit bumping rule, it allows traffic to be bumped, and it belongs to the protected group of the bundle. The QoS type of a VC using this class is ubr+.

```
vc-class atm priority-class
precedence 4-2
protect group
ubr+ 10000 3000
```

Basic-Class Class

When the class called basic-class is applied to a VC, the VC is configured through the **precedence other** command to carry traffic with IP Precedence levels not specified in the profile. The VC using this class belongs to the protected group of the bundle. The QoS type of a VC using this class is ubr.

```
vc-class atm basic-class
precedence other
protect group
ubr 10000
```

The following sets of commands configure three bundles that the router subinterface uses to connect to three of its neighbors. These bundles are called new-york, san-francisco, and los-angeles. Bundle new-york has four VC members, bundle san-francisco has four VC members, and bundle los-angeles has three VC members.

new-york Bundle

The first part of this example specifies the IP address of the subinterface, the router protocol--the router uses IS-IS as an IP routing protocol--and it creates the first bundle called new-york and enters bundle configuration mode:

```
interface atm 1/0.1 multipoint
ip address 10.0.0.1 255.255.255.0
ip router isis
bundle new-york
```

From within bundle configuration mode, the next portion of the configuration uses two protocol commands to enable IP and Open Systems Interconnect (OSI) traffic flows in the bundle. The OSI routing packets will use the highest precedence VC in the bundle. The OSI data packets, if any, will use the lowest precedence

VC in the bundle. If configured, other protocols, such as IPX or AppleTalk, will always use the lowest precedence VC in the bundle.

As the indentation levels of the preceding and following commands suggest, subordinate to bundle new-york is a command that configures its protocol and a command that applies the class called bundle-class to it.

```
protocol ip 1.1.1.2 broadcast
protocol clns 49.0000.0000.2222.00 broadcast
class-bundle bundle-class
```

The class called bundle-class, which is applied to the bundle new-york, includes a **protocol ip inarp** command. According to inheritance rules, **protocol ip**, configured at the bundle level, takes precedence over **protocol ip inarp** specified in the class bundle-class.

The next set of commands beginning with **pvc-bundle ny-control 207**, which are further subordinate, add four VCs (called ny-control, ny-premium, ny-priority, and ny-basic) to the bundle new-york. A particular class--that is, one of the classes predefined in this configuration example--is applied to each VC to configure it with parameters specified by commands included in the class.

As is the case for this configuration, to configure individual VCs belonging to a bundle, the router must be in bundle mode for the mother bundle. For each VC belonging to the bundle, the subordinate mode is pvc-mode for the specific VC.

The following commands configure the individual VCs for the bundle new-york:

```
pvc-bundle ny-control 207
  class-vc control-class
pvc-bundle ny-premium 206
  class-vc premium-class
pvc-bundle ny-priority 204
  class-vc priority-class
pvc-bundle ny-basic 201
  class-vc basic-class
```

san-francisco Bundle

The following set of commands create and configure a bundle called san-francisco. At the bundle configuration level, the configuration commands included in the class bundle-class are ascribed to the bundle san-francisco and to the individual VCs that belong to the bundle. Then, the **pvc-bundle** command is executed for each individual VC to add it to the bundle. After a VC is added and bundle-vc configuration mode is entered, a particular, preconfigured class is assigned to the VC. The configuration commands comprising that class are used to configure the VC. Rules of hierarchy apply at this point. Command parameters contained in the applied class are superseded by the same parameters applied at the bundle configuration level, which are superseded by the same parameters applied directly to a VC.

```
bundle san-francisco
protocol clns 49.0000.0000.0000.333.00 broadcast
inarp 1
class-bundle bundle-class
pvc-bundle sf-control 307
  class-vc control-class
pvc-bundle sf-premium 306
  class-vc premium-class
pvc-bundle sf-priority 304
  class-vc priority-class
pvc-bundle sf-basic 301
  class-vc basic-class
```

los-angeles Bundle

The following set of commands create and configure a bundle called los-angeles. At the bundle configuration level, the configuration commands included in the class bundle-class are ascribed to the bundle los-angeles and to the individual VCs that belong to the bundle. Then, the **pvc-bundle** command is executed for each individual VC to add it to the bundle. After a VC is added and bundle-vc configuration mode is entered, precedence is set for the VC and the VC is either configured as a member of a protected group (protect group) or as an individually protected VC. A particular class is then assigned to each VC to further characterize it. Rules of hierarchy apply. Parameters of commands applied directly and discretely to a VC take precedence over the same parameters applied within a class to the VC at the bundle-vc configuration level, which take precedence over the same parameters applied to the entire bundle at the bundle configuration level.

```
bundle los-angeles
  protocol ip 1.1.1.4 broadcast
  protocol clns 49.0000.0000.4444.00 broadcast
  inarp 1
  class-bundle bundle-class
  pvc-bundle la-high 407
    precedence 7-5
    protect vc
    class-vc premium-class
  pvc-bundle la-mid 404
    precedence 4-2
    protect group
    class-vc priority-class
  pvc-bundle la-low 401
    precedence other
    protect group
    class-vc basic-class
```

Example Per-VC WFQ and CBWFQ on a Standalone VC

The following example creates two class maps and defines their match criteria. For the first map class, called class1, the numbered access control list (ACL) 101 is used as the match criterion. For the second map class called class2, the numbered ACL 102 is used as the match criterion.

Next, the example includes these classes in a policy map called policy1. For class1, the policy includes a minimum bandwidth allocation request of 500 kbps and maximum packet count limit of 30 for the queue reserved for the class. For class2, the policy specifies only the minimum bandwidth allocation request of 1000 kbps, so the default queue limit of 64 packets is assumed. Note that the sum of the bandwidth requests for the two classes comprising policy1 is 75 percent of the total amount of bandwidth (2000 kbps) for the PVC called cisco to which the policy map is attached.

The example attaches the policy map called policy1 to the PVC called cisco. Once the policy map policy1 is attached to PVC cisco, its classes constitute the CBWFQ service policy for that PVC. Packets sent on this PVC will be checked for matching criteria against ACLs 101 and 102 and classified accordingly.

Because the **class-default** command is not explicitly configured for this policy map, all traffic that does not meet the match criteria of the two classes comprising the service policy is handled by the predefined class-default class, which provides best-effort flow-based WFQ.

```
class-map class1
  match access-group 101
class-map class2
  match access-group 102
policy-map policy1
  class class1
    bandwidth 500
    queue-limit 30
  class class2
    bandwidth 1000
interface ATM1/1/0.46 multipoint
  ip address 200.126.186.2 255.255.255.0
```

```
pvc cisco 46
  vbr-nrt 2000 2000
  encaps aal5snap
  service policy output policy1
```

Example Per-VC WFQ and CBWFQ on Bundle-Member VCs

The following example shows a PVC bundle called san-francisco with members for which per-VC WFQ and CBWFQ are enabled and service policies configured. The example assumes that the classes included in the following policy maps have been defined and that the policy maps have been created: policy1, policy2, and policy4. For each PVC, the IP to ATM CoS **pvc-bundle** command is used to specify the PVC to which the specified policy map is to be attached.

Note that PVC 0/34 and 0/31 have the same policy map attached to them, policy2. Although you can assign the same policy map to multiple VCs, each VC can have only one policy map attached at an output PVC.

```
bundle san-francisco
  protocol ip 1.0.2.20 broadcast
  encapsulation aal5snap
  pvc-bundle 0/35
    service policy output policy1
    vbr-nrt 5000 3000 500
    precedence 4-7
  pvc-bundle 0/34
    service policy output policy2
    vbr-nrt 5000 3000 500
    precedence 2-3
  pvc-bundle 0/33
    vbr-nrt 4000 3000 500
    precedence 2-3
    service policy output policy4
  pvc-bundle 0/31
    service policy output policy2
```

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned 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. (1005R)

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

