IP to ATM CoS

Feature Summary

The IP to ATM Class of Service (IP to ATM CoS) feature implements a solution for coarse-grained mapping of quality of service (QoS) characteristics between IP and ATM, using Cisco's PA-A3 ATM port adapters on 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 reinvestment of new ATM network infrastructures. Now networks can offer different service classes (sometimes termed "differential service classes") across the entire wide-area network (not just the routed portion). Mission-critical applications can be given exceptional service during periods of high network usage and congestion. In addition, non-critical traffic can be restricted in its network usage, which ensures greater quality of service for more important traffic and user types.

With this release of IP to ATM CoS (phase 1), network managers can use existing features (such as Committed Access Rate (CAR) or Policy-Based Routing) to classify and mark different IP traffic by modifying the IP Precedence field in the IPv4 packet header. Subsequently, Weighted Random Early Detection (WRED) can be configured on a per virtual circuit (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.

PA-A3 ATM port adapters provide the ability to shape traffic on each VC according to the ATM service category and traffic parameters employed. Using 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.

This release of IP to ATM CoS (phase 1) supports a single ATM VC, as shown in Figure 1. A future release of IP to ATM CoS (phase 2) will allow users to build a "VC bundle" that contains multiple VCs with different ATM QoS characteristics.

Figure 1 Single ATM Circuit Class

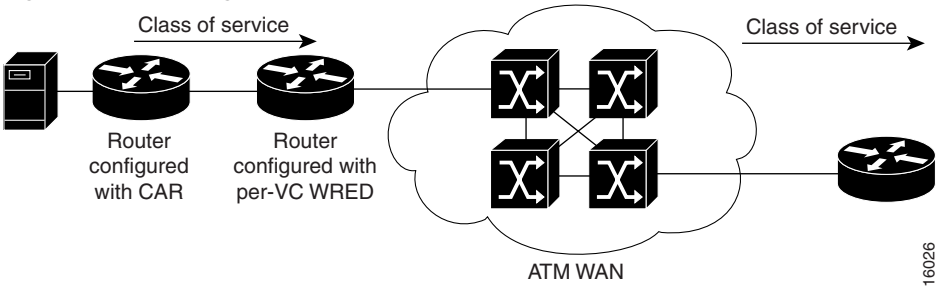

Benefits

IP to ATM Class of Service (IP to ATM CoS) provides the following benefits:

- Ensures effective differential classes over IP and traditional ATM networks
- Uses existing ATM infrastructures
- High performance design benefiting from distributed processing on the Cisco 7500 and VIP
- Per-VC queueing on the port adapter, per-VC back pressure and per-VC WRED VIP queueing bring stability to a network by ensuring that system packets (such as BGP and ISIS) are never dropped.

List of Terms

available bit rate (ABR)—An ATM service category in which the network may instruct sources to reduce their rate during times of congestion.

constant bit rate (CBR)—An ATM service category that aims to emulate a dedicated circuit of a certain fixed bandwidth.

IP Precedence—A 3-bit value in the Type of Service (ToS) byte of the IPv4 Header, which is used by WRED as a drop preference indicator.

permanent virtual circuit (PVC)—Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection.

port adapter—Media-specific interface PCI daughter card for use on the VIP.

Random Early Detection (RED)—An algorithm which, when applied, specifies that a small percentage of packets are to be dropped when congestion is detected—that is, before the queue in question overflows completely.

Type of Service (ToS)—A byte in IPv4 Header used, for example, by WRED as a drop preference indicator.

uncommitted bit rate (UBR)—An ATM service category defined by the ATM Forum for best effort traffic with no traffic-related service guarantees. No ATM traffic-related parameters are specified. A UBR circuit is by definition a best-effort circuit.

variable bit rate (VBR)—An ATM service category in which mean cell rate, peak cell rate, and burst tolerance are specified. A VBR circuit takes precedence over a UBR circuit in the event that there is contention for network resources.

virtual circuit (VC)—Logical circuit created to ensure reliable communication between two network devices. A virtual circuit (VC) is defined by a VPI/VCI pair, and can be either permanent (PVC) or switched (SVC). In ATM, a virtual circuit is called a virtual channel.

virtual channel identifier (VCI)— 16-bit field in the header of an ATM cell. The VCI, together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination.

Virtual Interface Processor (VIP)—Architecture for intelligent interface processors for the Cisco 7000 series routers. This architecture supports two port adapters, standard packet delivery, and distributed fast switching and feature offload.

virtual path identifier (VPI)—8-bit field in the header of an ATM cell. The VPI, together with the VCI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination.

Weighted Random Early Detection (WRED)—A variant of Random Early Detection (RED) in which the probability of a packet being dropped depends on its precedence, as well as other factors in the RED algorithm.

Weighted Fair Queueing (WFQ)—A queueing algorithm that provides a fraction of link bandwidth (constituting the weight) to each of several queues.

Restrictions

The following restrictions apply to using IP to ATM Class of Service (IP to ATM CoS):

- The ATM PA-A3 port adapter does not presently support available bit rate (ABR). IP to ATM CoS will support ABR when ABR is available on the ATM PA-A3 port adapter.
- The VIP2-50 supports only one ATM PA-A3 port adapter.
- This release of IP to ATM CoS (phase 1) accommodates only a single ATM circuit per next hop.
- IP to ATM CoS only supports permanent virtual circuits (PVCs).

Platforms

This feature is supported on any Cisco 7500 series router that is equipped with the following hardware:

- VIP2-50
- One of the following ATM port adapters:
  - PA-A3-T3
  - PA-A3-E3
  - PA-A3-OC3MM
  - PA-A3-OC3SMI
  - PA-A3-OC3SML

Prerequisites

To use this feature, distributed CEF must be enabled. You should also be familiar with the following Quality of Service (QoS) features:

- WRED
- CAR
- Policy-based routing

Documentation for these features can be found on the Documentation CD-ROM and on Cisco Connection Online (CCO).

Supported MIBs and RFCs

None.

For descriptions of supported MIBs and how to use MIBs, see Cisco's MIB website on CCO at http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml.

Functional Description

This section gives a broad overview of the IP to ATM Class of Service (IP to ATM CoS) feature. It includes the following topics:

- Why Use IP to ATM CoS
- IP to ATM CoS Features
- Congestion Avoidance

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.

Using Cisco's PA-A3 ATM port adapters on Cisco7500 series routers, the IP to ATM CoS feature solution provides the ability to map IP class of service and ATM quality of service, extending the capability previously available only for IP networks; differentiated services are preserved through the ATM network.

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 virtual circuit basis. Packets are queued and dequeued based on the back pressure from the port adapter. Having 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

  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 statistics

  This feature maintains per-flow and per-VC statistics based on IP Precedence.

Congestion Avoidance

For each VC that is created on the PA-A3 ATM port adapter, the port adapter allocates some of the buffers from its buffer pool to that VC in order to create a queue for that VC. Since the PA-A3 ATM port adapter has a finite buffer pool, the port adapter could run out of buffers when many VCs are configured.

To overcome this limitation, the IP to ATM CoS feature provides the ability to enable the buffering for a VC to be extended to the Layer 3 processor system, where more packets can be buffered for a particular VC. This is achieved by creation of queues in the Layer 3 processor system, which are then mapped one-to-one to the per-VC queues on the port adapter. When the port adapter 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, since the Layer 3 queues are accessible by the Layer 3 processor, a user can run flexible software scheduling algorithms on those queues.

When transporting 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 is reassembled and found to be incomplete.

To initiate effective congestion management techniques, IP to ATM CoS uses per virtual circuit-Weighted Random Early Detection (per-VC WRED). Per-VC WRED selectively places TCP sessions in slow start to ensure higher aggregate throughput under congestion. The example depicted in Figure 2 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 transmitted.

Figure 2 Traffic Congestion with IP to ATM CoS and Per-VC WRED

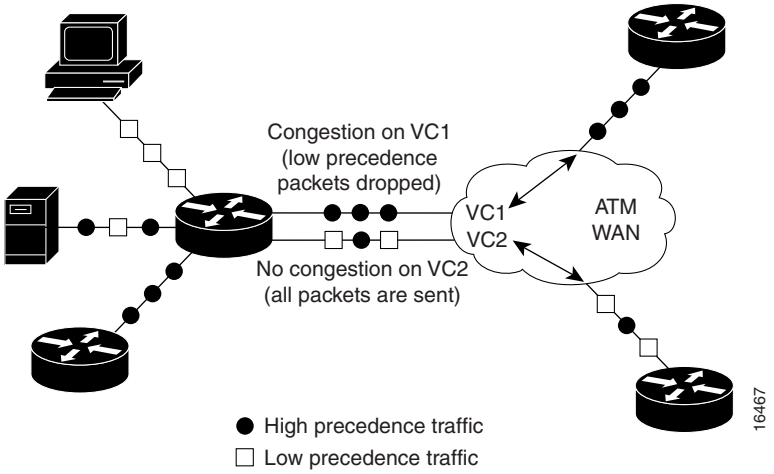

Running the WRED algorithm independently on each per-VC queue provides differentiated QoS to traffic of different IP Precedence values.

Configuration Tasks

To configure IP to ATM CoS, complete the following configuration tasks. The first task is required. The remaining are optional.

- Define the WRED Parameter Group
- Configure the Exponential Weight Factor (Optional)
- Configure the IP Precedence (Optional)
- View the WRED Parameters (Optional)
- View the Queueing Statistics (Optional)

Define the WRED Parameter Group

To define the Weighted Random Early Detection (WRED) parameter group, use the following command in global configuration mode:

Command	Purpose
random-detect-group *group-name*	Define the Weighted Random Early Detection (WRED) parameter group.

Configure the Exponential Weight Factor (Optional)

To configure the exponential weight factor for the average queue size calculation for a WRED parameter group, use the following commands beginning in global configuration mode:

Step	Command	Purpose
1	**random-detect-group** *group-name*	Specify the Weighted Random Early Detection (WRED) parameter group.
2	**exponential-weighting-constant** *exponent*	Configure the exponential weight factor for the average queue size calculation for the specified WRED parameter group.

Configure the IP Precedence (Optional)

To configure a WRED parameter group for a particular IP precedence, use the following commands beginning in global configuration mode:

Step	Command	Purpose
1	**random-detect-group** *group-name*	Specify the Weighted Random Early Detection (WRED) parameter group.
2	**precedence** *precedence min-threshold max-threshold mark-probability-denominator*	Configure the specified WRED parameter group for a particular IP Precedence.

View the WRED Parameters (Optional)

To view the configured WRED parameters, use the following command beginning in global configuration mode:

Command	Purpose
show queueing red [**interface** *atm_subinterface* [**vc**[[*vpi*/] *vci*]]]	Show the parameters of every VC with WRED enabled on the specified ATM subinterface.

View the Queueing Statistics (Optional)

To view the queueing statistics for an interface, use the following command in global configuration mode:

Command	Purpose
show queueing interface *interface-number* [**vc** [[*vpi*/] *vci*]]	Show the queueing statistics of a specific VC on an interface.

Configuration Example

The following example creates a PVC on an ATM interface and applies the WRED parameter group "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
 atm pvc 46 0 46 aal5nlpid random-detect 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
!
```

Command Reference

This section documents new or modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.0 command references.

- **atm pvc**
- **exponential-weighting-constant**
- **precedence**
- **random-detect-group**
- **show queueing interface**
- **show queueing red**

atm pvc

To create a permanent virtual circuit (PVC) on an ATM interface and, optionally, to generate Operation, Administration, and Maintenance (OAM) F5 loopback cells or enable Inverse ATM ARP, use the **atm pvc** interface configuration command. To apply a WRED parameter group to the created PVC, specify the **random-detect** keyword and *group-name* argument. The **no** form of this command removes the specified PVC.

atm pvc *vcd vpi vci aal-encap* [[*midlow midhigh*] [*peak average* [*burst*]]] [**inarp** [*minutes*]] [**oam** [*seconds*] [**random-detect** [*group-name*]]
no atm pvc *vcd vpi vci aal-encap* [[*midlow midhigh*] [*peak average* [*burst*]]] [**inarp** [*minutes*]] [**oam** [*seconds*][**random-detect** [*group-name*]]

Syntax Description

vcd — Virtual circuit descriptor. A unique number that identifies to the processor which VPI-VCI pair to use for a particular packet. Valid values range from 1 to the value set with the **atm maxvc** command. The AIP or ATM port adapter requires this feature to manage packet transmission. The *vcd* value is not associated with the VPI-VCI pair used for the ATM network cells. The NPM has a hard coded max *vcd* value of 1023.

vpi — ATM network virtual path identifier (VPI) of this PVC. On the Cisco 7200 and 7500 series routers, this value ranges from 0 to 255; on the Cisco 4500 and Cisco 4700 series, this value ranges from 0 to 1 less than the quotient of 8192 divided by the value set by the **atm vc-per-vp** command.

The VPI is an 8-bit field in the header of the ATM cell. The VPI value is unique only on a single link, not throughout the ATM network, because it has local significance only. The VPI value must match that of the switch.

The arguments *vpi* and *vci* cannot both be set to 0; if one is 0, the other cannot be 0.

vci

ATM network virtual channel identifier (VCI) of this PVC, in the range of 0 to 1 less than the maximum value set for this interface by the **atm vc-per-vp** command. Typically, lower values 0 to 31 are reserved for specific traffic (for example, F4 OAM, SVC signaling, ILMI, and so on) and should not be used.

The VCI is a 16-bit field in the header of the ATM cell. The VCI value is unique only on a single link, not throughout the ATM network, because it has local significance only.

The arguments *vpi* and *vci* cannot both be set to 0; if one is 0, the other cannot be 0.

aal-encap

ATM adaptation layer (AAL) and encapsulation type. When the **aal5mux** keyword is specified, a protocol is required. Possible values are as follows:

- **aal34smds**—Encapsulation for SMDS networks. This option is supported on the AIP and is not available for the ATM port adapter.
- **aal5nlpid**—Encapsulation that allows ATM interfaces to interoperate with High-Speed Serial Interfaces (HSSIs) that are using an ATM data service unit (ADSU) and running ATM-Data Exchange Interface (DXI).
- **aal5mux apollo**—A multiplex (MUX)-type virtual circuit.
- **aal5mux appletalk**—A MUX-type virtual circuit.
- **aal5mux decnet**—A MUX-type virtual circuit.
- **aal5mux ip**—A MUX-type virtual circuit.
- **aal5mux ipx**—A MUX-type virtual circuit.
- **aal5mux vines**—A MUX-type virtual circuit.
- **aal5mux xns**—A MUX-type virtual circuit.
- **aal5snap**—Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) precedes the protocol datagram. This is the only encapsulation supported for Inverse ARP.
- **ilmi**—Used to set up communication with the ILMI; the associated *vpi* and *vci* values are ordinarily 0 and 16, respectively.
- **qsaal**—A signaling-type PVC used for setting up or tearing down SVCs; the associated *vpi* and *vci* values are ordinarily 0 and 5, respectively.

midlow

(Set for the **aal34smds** encapsulation only) (Optional) Starting message identifier (MID) number for this PVC. The default is 0. If you set the *peak* and *average* (*burst* is optional) values for **aal34smds** encapsulation, you must also set the *midlow* and *midhigh* values.

This option is not available for the ATM port adapter.

midhigh

(Set for the **aal34smds** encapsulation only) (Optional) Ending MID number for this PVC. The default is 0. If you set the *peak* and *average* (*burst* is optional) values for **aal34smds** encapsulation, you must also set the *midlow* and *midhigh* values.

This option is not available for the ATM port adapter.

peak	(Optional) Maximum rate (in kbps) at which this virtual circuit can transmit. Valid values are in the range from 1 to the maximum rate set for a rate queue. If you set this value, you must also specify the *average* (*burst* is optional) value. If you set the *peak* and *average* values for **aal34smds** encapsulation, you must also set the *midlow* and *midhigh* values. This option is not available for the ATM port adapter. It is available on the Enhanced ATM port adapter (PA-A3).
average	(Optional) Average rate (in kbps) at which this virtual circuit transmits. Valid values are platform dependent. If you set this value, you must also specify the *peak* (*burst* is optional) value. If you set the *peak* and *average* values for **aal34smds** encapsulation, you must also set the *midlow* and *midhigh* values. This option is not available for the ATM port adapter. It is available on the Enhanced ATM port adapter (PA-A3).
burst	(Optional) Value that relates to the maximum number of ATM cells the virtual circuit can transmit to the network at the *peak* rate of the PVC. On the AIP, the actual burst cells equals *burst* * 32 cells, thereby allowing for a burst size of 32 cells to 2016 cells. On the ATMZR the value is not multiplied. If you set this value, you must also specify a value for the *peak* and *average* values. On the AIP, *burst* can range from 1 to 63 On the ATMZR, *burst* can range from 1 to 65535 This option is not available for the ATM port adapter. It is available on the Enhanced ATM port adapter (PA-A3).
inarp *minutes*	(Set for the **aal5snap** encapsulation only) (Optional) Specifies how often inverse ARP datagrams are sent on this virtual circuit. The default value is 15 minutes.
oam *seconds*	(Optional) Specifies how often to generate an OAM F5 loopback cell from this virtual circuit. The default value is 10 seconds.
random-detect *group-name*	(Set for the IP to ATM Cos feature only.) (Optional) Specifies that the WRED algorithm should be applied to this PVC. If the *group-name* argument is omitted or no name match is found in the current configuration for the specified *group-name,* a set of default WRED parameters are applied to the PVC. Configuring a PVC without the **random-detect** keyword turns WRED off for the PVC.

Defaults

If *peak* and *average* rate values are omitted, the PVC defaults to peak and average rates equal to the link rate. The peak and average rates are then equal. By default, the virtual circuit is configured to run as fast as possible.

The default of both the *midlow* and *midhigh* values is 0.

If the **oam** keyword is omitted, OAM cells are not generated. If the **oam** keyword is present but the *seconds* value is omitted, the default value of **oam** *seconds* is 10 seconds.

If the **inarp** keyword is omitted, inverse ARPs are not generated. If the **inarp** keyword is present, but the timeout value is not given, then inverse ARPs are generated every 15 minutes.

When the **random-detect** keyword is used, if the *group-name* argument is omitted or no name match is found in the current configuration for the specified *group-name*, a set of default WRED parameters are applied to the PVC.

If the **random-detect** keyword is omitted, WRED is turned off for the PVC.

Command Mode

Interface configuration

Usage Guidelines

This command first appeared in Cisco IOS Release 10.0. The *midlow* and *midhigh* arguments first appeared in Cisco IOS Release 10.3. The **oam** *seconds* and **inarp** *minutes* arguments first appeared in Cisco IOS Release 11.0. The **random-detect** keyword first appeared in Cisco IOS Release 11.1(22)CC.

Because the ATM port adapters do not support traffic shaping, the *peak*, *average*, and *burst* rate options are not available. For more information on the ATM port adapter, refer to the *PA-A1 ATM Port Adapter Installation and Configuration* publication.

The order of command options is important. The **inarp** keyword can be specified either separately or before the **oam** keyword has been enabled. The *peak*, *average*, and *burst* arguments, if specified, cannot be specified after either the **inarp** or the **oam** keywords.

The Cisco IOS software dynamically creates rate queues as necessary to satisfy the requests of **atm pvc** commands. The software dynamically creates a rate queue when an **atm pvc** command specifies a peak or average rate that does not match any user-configured rate queue.

The **atm pvc** command creates a PVC and attaches it to the VPI and VCI specified. Both *vpi* and *vci* cannot be specified as 0; if one is 0, the other cannot be 0. The *aal-encap* argument determines the AAL mode and the encapsulation method used. The *peak* and *average* arguments determine the rate queue used.

Use one of the **aal5mux** encapsulation options to dedicate the specified virtual circuit to a single protocol; use the **aal5snap** encapsulation option to multiplex two or more protocols over the same virtual circuit. Whether you select **aal5mux** or **aal5snap** encapsulation might depend on practical considerations, such as the type of network and the pricing offered by the network. If the network's pricing depends on the number of virtual circuits set up, **aal5snap** might be the appropriate choice. If pricing depends on the number of bytes transmitted, **aal5mux** might be the appropriate choice because it has slightly less overhead.

If you choose to specify *peak* or *average* values, you must specify both. If you set the *peak* and *average* values for **aal34smds** encapsulation, you must also specify the *midlow* and *midhigh* values. **aal34smds** encapsulation is not available for the ATM port adapter.

In the 7000 router series family (AIP), the values for *peak* and *average* indicate the bandwidth as seen on the wire.

In the 4500 router series family, (ATMizer), the values for *peak* and *average* indicate the bandwidth of the AAL5 payload (exclusive of padding).

Message identifier (MID) numbers, which are available only with **aal34smds** encapsulation, are used by receiving devices to reassemble cells from multiple packets. You can assign different *midlow* to *midhigh* ranges to different PVCs to ensure that the message identifiers are unique at the receiving end and, therefore, that messages can be reassembled correctly.

When configuring an SVC, use the **atm pvc** command to configure the PVC that handles the SVC call setup and termination. In this case, specify the **qsaal** encapsulation for the *aal-encap* keyword. See the third example that follows.

The router generates and echoes OAM F5 loopback cells, which verify connectivity. Once OAM cell generation is enabled, a cell is transmitted periodically. The remote end must respond by echoing back the cells.

The router does not generate alarm indication signal (AIS) cells, which are used for alarm surveillance functions. However, if it receives an AIS cell, it responds by sending an OAM far-end remote failure (FERF) cell.

Examples

The following example creates a PVC with VPI 0 and VCI 6. The PVC uses AAL AAL5-MUX with IP protocol.

```
atm pvc 1 0 6 aal5mux ip
```

The following example creates a PVC with VPI 0 and VCI 6. The PVC uses AAL AAL3/4-SMDS protocol.

```
atm pvc 1 0 6 aal34smds 0 15 150000 70000 10
```

The following example creates a PVC to be used for ATM signaling for an SVC. It specifies VPI 0 and VCI 5.

```
atm pvc 1 0 5 qsaal
```

Assuming that no static rate queue has been defined, the following example creates the PVC and also creates a dynamic rate queue with the peak rate set to the maximum allowed by the physical layer interface module (PLIM) and the average set to equal the peak rate:

```
atm pvc 1 1 1 aal5snap
```

Assuming that no static rate queue has been defined, the following example creates the PVC and also creates a dynamic rate queue with the peak rate set to 100 Mbps (100,000 kbps), the average rate set to 50 Mbps (50,000 kbps), and a burst size of 64 cells (2 * 32 cells):

```
atm pvc 1 1 1 aal5snap 100000 50000 2
```

The following example creates a PVC to be used for IP to ATM Class of Service (IP to ATM CoS). It specifies VPI 0 and VCI 46. The PVC uses aal5nlpid encapsulation and the WRED parameter group “sanjose” is applied to the PVC.

```
atm pvc 1 0 46 aal5nlpid random-detect sanjose
```

Related Commands

atm aal aal3/4
atm maxvc
atm multicast
atm rate-queue
atm smds-address
mtu

random-detect-group
show queueing interface
show queueing red

exponential-weighting-constant

To configure the exponential weight factor for the average queue size calculation for a WRED parameter group, use the **exponential-weighting-constant** command. The **no** form of this command returns the group's value to the default.

exponential-weighting-constant *exponent*
no exponential-weighting-constant

Syntax Description

exponent	Exponent from 1 to 16 used in the average queue size calculation. The default is 9.

Default

The weight factor is 9.

Command Mode

random-detect-group

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1(22) CC. If used, it is issued after the **random-detect-group** command.

Use this command to change the exponent used in the average queue size calculation for a WRED parameter group. The average queue size is based on the previous average and the current size of the queue. The formula is

```
average = (old_average * (1-1/2^n)) + (current_queue_size * 1/2^n)
```

where *n* is the exponential weight factor specified in this command. Thus, the higher the factor, the more dependent the average is on the previous average.

Note The default WRED parameter values are based on the best available data. Cisco recommends that you do not change the parameters from their default values unless you have determined that your applications would benefit from the changed values.

For high values of *n*, the previous average becomes more important. A large factor smooths out the peaks and lows in queue length. The average queue size is unlikely to change very quickly, avoiding drastic swings in size. The WRED process will be slow to start dropping packets, but it may continue dropping packets for a time after the actual queue size has fallen below the minimum threshold. The slow-moving average will accommodate temporary bursts in traffic.

If the value of *n* gets too high, WRED will not react to congestion. Packets will be transmitted or dropped as if WRED were not in effect.

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

If the value of *n* gets too low, WRED will overreact to temporary traffic bursts and drop traffic unnecessarily.

Example

The following example configures the WRED group "sanjose" with a weight factor of 10:

```
random-detect-group sanjose
  exponential-weighting-constant 10
```

Related Commands

precedence
random-detect exponential-weighting constant
random-detect-group
show queueing interface
show queueing red

precedence

To configure a WRED parameter group for a particular IP Precedence, use the **precedence** command. The **no** form of this command returns the group's values to the default for the IP Precedence.

precedence *precedence min-threshold max-threshold mark-probability-denominator*
no precedence *precedence min-threshold max-threshold mark-probability-denominator*

Syntax Description

precedence	IP precedence number. The value ranges from 0 to 7.
min-threshold	Minimum threshold in number of packets. The value ranges from 1 to 4096. When the average queue length reaches this number, RED begins to drop packets with the specified IP Precedence.
max-threshold	Maximum threshold in number of packets. The value ranges from *min-threshold* to 4096. When the average queue length exceeds this number, WRED drops all packets with the specified IP Precedence.
mark-probability-denominator	Denominator for the fraction of packets dropped when the average queue depth is *max-threshold*. For example, if the denominator is 512, one out of every 512 packets is dropped when the average queue is at the *max-threshold*. The value is from 1 to 65536. The default is 10; one out of every ten packets are dropped at the *max-threshold*.

Default

For all IP Precedences, the *mark-probability-denominator* is 10, and the *max-threshold* is based on the output buffering capacity and the transmission speed for the interface.

The default *min-threshold* depends on the IP Precedence. The *min-threshold* for IP Precedence 0 corresponds to one half of the *max-threshold*. The values for the remaining IP Precedences fall between one half the *max-threshold* and the *max-threshold* at evenly spaced intervals. Table 1 lists the default minimum value for each IP Precedence.

Table 1 Default WRED Minimum Threshold Values

IP Precedence	Minimum Threshold Value (Fraction of Maximum Threshold Value)
0	8/16
1	9/16
2	10/16
3	11/16
4	12/16
5	13/16

Table 1 Default WRED Minimum Threshold Values (continued)

IP Precedence	Minimum Threshold Value (Fraction of Maximum Threshold Value)
6	14/16
7	15/16

Command Mode

random-detect-group

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1(22)CC. If used, it is issued after the **random-detect-group** command.

When you configure the **random-detect group** command on an interface, packets are given preferential treatment based on the IP Precedence of the packet. Use the **precedence** command to adjust the treatment for different IP Precedences.

If you want WRED to ignore the IP Precedence when determining which packets to drop, enter this command with the same parameters for each IP Precedence. Remember to use reasonable values for the minimum and maximum thresholds.

Note The default WRED parameter values are based on the best available data. We recommend that you do not change the parameters from their default values unless you have determined that your applications would benefit from the changed values.

Example

The following example specifies parameters for WRED parameter group "sanjose" for the different IP Precedences:

```
random-detect-group sanjose
  precedence 0 32 256 100
  precedence 1 64 256 100
  precedence 2 96 256 100
  precedence 3 128 256 100
  precedence 4 160 256 100
  precedence 5 192 256 100
  precedence 6 224 256 100
  precedence 7 256 256 100
```

Related Commands

exponential-weighting-constant
random-detect precedence
random-detect-group
show queueing interface
show queueing red

random-detect-group

To define the Weighted Random Early Detection (WRED) parameter group, use the **random-detect group** global configuration command. The **no** form of this command deletes the WRED parameter group.

random-detect-group *group-name*
no random-detect-group *group-name*

Syntax Description

group-name	Name for the WRED parameter group.

Default

No WRED parameter group exists.

Command Mode

Global configuration

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1(22) CC.

WRED is a congestion avoidance mechanism that slows traffic by randomly dropping packets when there is congestion. WRED is most useful when the traffic uses protocols such as TCP, which respond to dropped packets by decreasing the transmission rate.

The router automatically determines parameters to use in the WRED calculations. If you want to change these parameters for a group, use the **exponential-weighting-constant** or **precedence** commands.

Example

The following example defines the WRED parameter group "sanjose:"

```
random-detect-group sanjose
```

Related Commands

atm pvc
exponential-weighting-constant
precedence
random detect
show queueing interface
show queueing red

show queueing interface

To show the queueing statistics of an interface, use the **show queueing interface** privileged EXEC command.

show queueing interface *interface-number* [**vc** [[*vpi***/**] *vci*]]

Syntax Description

interface-number	Specifies the interface.
vc	(Optional) Shows the WRED parameters associated with a specific VC. If desired, both the VPI and VCI values, or just the VCI value, can be specified.
*vpi***/**	(Optional) Specifies the virtual circuit identifier (VPI). If *vpi* is omitted, 0 is used as the VPI value for locating the PVC. If *vpi* is specified, the **/** separator is required.
vci	(Optional) Specifies the virtual path identifier (VCI).

Command Mode

Privileged EXEC

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1(22)CC.

Related Commands

atm pvc
exponential-weighting-constant
precedence
random-detect-group
show queueing
show queueing red

show queueing red

To display the configured WRED parameters, use the **show queueing red** privileged EXEC command.

show queueing red [**interface** *atm_subinterface* [**vc**[[*vpi***/**] *vci*]]]

Syntax Description

interface *atm_subinterface*	(Optional) Shows the WRED parameters of every VC with WRED enabled on the specified ATM subinterface.
vc	(Optional) Shows the WRED parameters associated with a specific VC. If desired, both the VPI and VCI values, or just the VCI value, can be specified.
*vpi***/**	(Optional) Specifies the virtual circuit identifier (VPI). If *vpi* is omitted, 0 is used as the VPI value for locating the PVC. If *vpi* is specified, the **/** separator is required.
vci	(Optional) Specifies the virtual path identifier (VCI).

Command Mode

Privileged EXEC

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1(22)CC

If no keyword is entered, this command displays the parameters of every WRED-enabled VC on every ATM interface in the system.

Related Commands

atm pvc
exponential-weighting-constant
precedence
random-detect-group
show queueing
show queueing interface