QoS: Policing and Shaping Configuration Guide, Cisco IOS Release 12.2SR

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

Policing and Shaping Overview

Cisco IOS QoS offers two kinds of traffic regulation mechanisms--policing and shaping.

The rate-limiting features of committed access rate (CAR) and the Traffic Policing feature provide the functionality for policing traffic. The features of Generic Traffic Shaping (GTS), Class-Based Traffic Shaping, Distributed Traffic Shaping (DTS), and Frame Relay Traffic Shaping (FRTS) provide the functionality for shaping traffic.

You can deploy these features throughout your network to ensure that a packet, or data source, adheres to a stipulated contract and to determine the QoS to render the packet. Both policing and shaping mechanisms use the traffic descriptor for a packet--indicated by the classification of the packet--to ensure adherence and service.

Policers and shapers usually identify traffic descriptor violations in an identical manner. They usually differ, however, in the way they respond to violations, for example:

- A policer typically drops traffic. (For example, the CAR rate-limiting policer will either drop the packet or rewrite its IP precedence, resetting the type of service bits in the packet header.)
- A shaper typically delays excess traffic using a buffer, or queueing mechanism, to hold packets and shape the flow when the data rate of the source is higher than expected. (For example, GTS and Class-Based Shaping use a weighted fair queue to delay packets in order to shape the flow, and DTS and FRTS use either a priority queue, a custom queue, or a FIFO queue for the same, depending on how you configure it.)

Traffic shaping and policing can work in tandem. For example, a good traffic shaping scheme should make it easy for nodes inside the network to detect misbehaving flows. This activity is sometimes called policing the traffic of the flow.

This module gives a brief description of the Cisco IOS QoS traffic policing and shaping mechanisms. Because policing and shaping all use the token bucket mechanism, this module first explains how a token bucket works. This module includes the following sections:

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.

What Is a Token Bucket

A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, a mean rate, and a time interval (Tc). Although the mean rate is generally represented as bits per second, any two values may be derived from the third by the relation shown as follows:

```
mean rate = burst size / time interval
```

Here are some definitions of these terms:

- Mean rate--Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average.
- Burst size--Also called the Committed Burst (Bc) size, it specifies in bits (or bytes) per burst, how much traffic can be sent within a given unit of time to not create scheduling concerns. (For a shaper, such as GTS, it specifies bits per burst; for a policer, such as CAR, it specifies bytes per burst, per second.)
- Time interval--Also called the measurement interval, it specifies the time quantum in seconds per burst.

By definition, over any integral multiple of the interval, the bit rate of the interface will not exceed the mean rate. The bit rate, however, may be arbitrarily fast within the interval.

A token bucket is used to manage a device that regulates the data in a flow. For example, the regulator might be a traffic policer, such as CAR, or a traffic shaper, such as FRTS or GTS. A token bucket itself has no discard or priority policy. Rather, a token bucket discards tokens and leaves to the flow the problem of managing its transmission queue if the flow overdrives the regulator. (Neither CAR nor FRTS and GTS implement either a true token bucket or true leaky bucket.)

In the token bucket metaphor, tokens are put into the bucket at a certain rate. The bucket itself has a specified capacity. If the bucket fills to capacity, newly arriving tokens are discarded. Each token is permission for the source to send a certain number of bits into the network. To send a packet, the regulator must remove from the bucket a number of tokens equal in representation to the packet size.

If not enough tokens are in the bucket to send a packet, the packet either waits until the bucket has enough tokens (in the case of GTS) or the packet is discarded or marked down (in the case of CAR). If the bucket is already full of tokens, incoming tokens overflow and are not available to future packets. Thus, at any time, the largest burst a source can send into the network is roughly proportional to the size of the bucket.

Note that the token bucket mechanism used for traffic shaping has both a token bucket and a data buffer, or queue; if it did not have a data buffer, it would be a policer. For traffic shaping, packets that arrive that cannot be sent immediately are delayed in the data buffer.

For traffic shaping, a token bucket permits burstiness but bounds it. It guarantees that the burstiness is bounded so that the flow will never send faster than the token bucket's capacity, divided by the time interval, plus the established rate at which tokens are placed in the token bucket. See the following formula:

```
(token bucket capacity in bits / time interval in seconds) + established rate in bps =
maximum flow speed in bps
```

This method of bounding burstiness also guarantees that the long-term transmission rate will not exceed the established rate at which tokens are placed in the bucket.

Policing with CAR

Committed access rate (CAR) embodies a rate-limiting feature for policing traffic, in addition to its packet classification feature discussed in the "Classification Overview" module. The rate-limiting feature of CAR manages the access bandwidth policy for a network by ensuring that traffic falling within specified rate parameters is sent, while dropping packets that exceed the acceptable amount of traffic or sending them with a different priority. The exceed action for CAR is to drop or mark down packets.

The rate-limiting function of CAR does the following:

- Allows you to control the maximum rate of traffic sent or received on an interface.
- Gives you the ability to define Layer 3 aggregate or granular incoming or outgoing (ingress or egress) bandwidth rate limits and to specify traffic handling policies when the traffic either conforms to or exceeds the specified rate limits.

Aggregate bandwidth rate limits match all of the packets on an interface or subinterface. Granular bandwidth rate limits match a particular type of traffic based on precedence, MAC address, or other parameters.

CAR is often configured on interfaces at the edge of a network to limit traffic into or out of the network.

- How CAR Works, page 3
- Restrictions of CAR and VIP-Distributed CAR, page 7

How CAR Works

CAR examines traffic received on an interface or a subset of that traffic selected by access list criteria. It then compares the rate of the traffic to a configured token bucket and takes action based on the result. For example, CAR will drop the packet or rewrite the IP precedence by resetting the type of service (ToS) bits. You can configure CAR to send, drop, or set precedence.

CAR utilizes a token bucket measurement. Tokens are inserted into the bucket at the committed rate. The depth of the bucket is the burst size. Traffic arriving at the bucket when sufficient tokens are available is said to conform, and the corresponding number of tokens are removed from the bucket. If a sufficient number of tokens are not available, then the traffic is said to exceed.

- Matching Criteria, page 3
- Rate Limits, page 4
- Conform and Exceed Actions, page 6
- Multiple Rate Policies, page 6

Matching Criteria

Traffic matching entails identification of traffic of interest for rate limiting, precedence setting, or both. Rate policies can be associated with one of the following qualities:

- Incoming interface
- All IP traffic
- IP precedence (defined by a rate-limit access list)
- MAC address (defined by a rate-limit access list)
- Multiprotocol Label Switching (MPLS) experimental (EXP) value (defined by a rate-limit access list)
- IP access list (standard and extended)

CAR provides configurable actions, such as send, drop, or set precedence when traffic conforms to or exceeds the rate limit.

Note Matching to IP access lists is more processor intensive than matching based on other criteria.

Rate Limits

CAR propagates bursts. It does no smoothing or shaping of traffic, and therefore does no buffering and adds no delay. CAR is highly optimized to run on high-speed links--DS3, for example--in distributed mode on Versatile Interface Processors (VIPs) on the Cisco 7500 series.

CAR rate limits may be implemented either on input or output interfaces or subinterfaces including Frame Relay and ATM subinterfaces.

What Rate Limits Define

Rate limits define which packets conform to or exceed the defined rate based on the following three parameters:

- Average rate. The average rate determines the long-term average transmission rate. Traffic that falls under this rate will always conform.
- Normal burst size. The normal burst size determines how large traffic bursts can be before some traffic exceeds the rate limit.
- Excess Burst size. The Excess Burst (Be) size determines how large traffic bursts can be before all traffic exceeds the rate limit. Traffic that falls between the normal burst size and the Excess Burst size exceeds the rate limit with a probability that increases as the burst size increases.

The maximum number of tokens that a bucket can contain is determined by the normal burst size configured for the token bucket.

When the CAR rate limit is applied to a packet, CAR removes from the bucket tokens that are equivalent in number to the byte size of the packet. If a packet arrives and the byte size of the packet is greater than the number of tokens available in the standard token bucket, extended burst capability is engaged if it is configured.

Extended Burst Value

Extended burst is configured by setting the extended burst value greater than the normal burst value. Setting the extended burst value equal to the normal burst value excludes the extended burst capability. If extended burst is not configured, given the example scenario, the exceed action of CAR takes effect because a sufficient number of tokens are not available.

When extended burst is configured and this scenario occurs, the flow is allowed to borrow the needed tokens to allow the packet to be sent. This capability exists so as to avoid tail-drop behavior, and, instead, engage behavior like that of Random Early Detection (RED).

How Extended Burst Capability Works

Here is how the extended burst capability works. If a packet arrives and needs to borrow *n* number of tokens because the token bucket contains fewer tokens than its packet size requires, then CAR compares the following two values:

- Extended burst parameter value.
- Compounded debt. Compounded debt is computed as the sum over all *ai*:
  - *a* indicates the actual debt value of the flow after packet *i* is sent. Actual debt is simply a count of how many tokens the flow has currently borrowed.
  - *i* indicates the *i*th packet that attempts to borrow tokens since the last time a packet was dropped.

If the compounded debt is greater than the extended burst value, the exceed action of CAR takes effect. After a packet is dropped, the compounded debt is effectively set to 0. CAR will compute a new compounded debt value equal to the actual debt for the next packet that needs to borrow tokens.

If the actual debt is greater than the extended limit, all packets will be dropped until the actual debt is reduced through accumulation of tokens in the token bucket.

Dropped packets do not count against any rate or burst limit. That is, when a packet is dropped, no tokens are removed from the token bucket.

Note Though it is true the entire compounded debt is forgiven when a packet is dropped, the actual debt is not forgiven, and the next packet to arrive to insufficient tokens is immediately assigned a new compounded debt value equal to the current actual debt. In this way, actual debt can continue to grow until it is so large that no compounding is needed to cause a packet to be dropped. In effect, at this time, the compounded debt is not really forgiven. This scenario would lead to excessive drops on streams that continually exceed normal burst. (See the example in the following section, "Actual and Compounded Debt Example, page 5."

Testing of TCP traffic suggests that the chosen normal and extended burst values should be on the order of several seconds worth of traffic at the configured average rate. That is, if the average rate is 10 Mbps, then a normal burst size of 10 to 20 Mb and an Excess Burst size of 20 to 40 Mb would be appropriate.

Recommended Burst Values

Cisco recommends the following values for the normal and extended burst parameters:

```
normal burst = configured rate * (1 byte)/(8 bits) * 1.5 seconds
extended burst = 2 * normal burst
```

With the listed choices for parameters, extensive test results have shown CAR to achieve the configured rate. If the burst values are too low, then the achieved rate is often much lower than the configured rate.

Actual and Compounded Debt Example

This example shows how the compounded debt is forgiven, but the actual debt accumulates.

For this example, assume the following parameters:

- Token rate is 1 data unit per time unit
- Normal burst size is 2 data units
- Extended burst size is 4 data units
- 2 data units arrive per time unit

After 2 time units, the stream has used up its normal burst and must begin borrowing one data unit per time unit, beginning at time unit 3:

```
Time    DU arrivals     Actual Debt     Compounded Debt
-------------------------------------------------------
1       2               0               0
2       2               0               0
3       2               1               1
4       2               2               3
5       2               3 (temporary)   6 (temporary)
```

At this time a packet is dropped because the new compounded debt (6) would exceed the extended burst limit (4). When the packet is dropped, the compounded debt effectively becomes 0, and the actual debt is 2. (The values 3 and 6 were only temporary and do not remain valid in the case where a packet is dropped.) The final values for time unit 5 follow. The stream begins borrowing again at time unit 6.

```
Time    DU arrivals     Actual Debt     Compounded Debt
-------------------------------------------------------
5       2               2               0
6       2               3               3
7       2               4 (temporary)   7 (temporary)
```

At time unit 6, another packet is dropped and the debt values are adjusted accordingly.

```
Time    DU arrivals     Actual Debt     Compounded Debt
-------------------------------------------------------
7       2               3               0
```

Conform and Exceed Actions

CAR utilizes a token bucket, thus CAR can pass temporary bursts that exceed the rate limit as long as tokens are available.

Once a packet has been classified as conforming to or exceeding a particular rate limit, the router performs one of the following actions on the packet:

- Transmit--The packet is sent.
- Drop--The packet is discarded.
- Set precedence and transmit--The IP Precedence (ToS) bits in the packet header are rewritten. The packet is then sent. You can use this action to either color (set precedence) or recolor (modify existing packet precedence) the packet.
- Continue--The packet is evaluated using the next rate policy in a chain of rate limits. If there is not another rate policy, the packet is sent.
- Set precedence and continue--Set the IP Precedence bits to a specified value and then evaluate the next rate policy in the chain of rate limits.

For VIP-based platforms, two more actions are possible:

- Set QoS group and transmit--The packet is assigned to a QoS group and sent.
- Set QoS group and continue--The packet is assigned to a QoS group and then evaluated using the next rate policy. If there is not another rate policy, the packet is sent.

Multiple Rate Policies

A single CAR rate policy includes information about the rate limit, conform actions, and exceed actions. Each interface can have multiple CAR rate policies corresponding to different types of traffic. For example, low priority traffic may be limited to a lower rate than high priority traffic. When there are multiple rate

policies, the router examines each policy in the order entered until the packet matches. If no match is found, the default action is to send.

Rate policies can be independent: each rate policy deals with a different type of traffic. Alternatively, rate policies can be cascading: a packet may be compared to multiple different rate policies in succession.

Cascading of rate policies allows a series of rate limits to be applied to packets to specify more granular policies (for example, you could rate limit total traffic on an access link to a specified subrate bandwidth and then rate limit World Wide Web traffic on the same link to a given proportion of the subrate limit) or to match packets against an ordered sequence of policies until an applicable rate limit is encountered (for example, rate limiting several MAC addresses with different bandwidth allocations at an exchange point). You can configure up to a 100 rate policies on a subinterface.

Restrictions of CAR and VIP-Distributed CAR

CAR and VIP-distributed CAR can only be used with IP traffic. Non-IP traffic is not rate limited.

CAR or VIP-distributed CAR can be configured on an interface or subinterface. However, CAR and VIP-distributed CAR are not supported on the following interfaces:

- Fast EtherChannel
- Tunnel
- PRI
- Any interface that does not support Cisco Express Forwarding (CEF)

CAR is only supported on ATM subinterfaces with the following encapsulations: aal5snap, aal5mux, and aal5nlpid.

Note CAR provides rate limiting and does not guarantee bandwidth. CAR should be used with other QoS features, such as distributed weighted fair queueing (DWFQ), if premium bandwidth assurances are required.

Traffic Policing

Traffic policing allows you to control the maximum rate of traffic sent or received on an interface, and to partition a network into multiple priority levels or class of service (CoS).

The Traffic Policing feature manages the maximum rate of traffic through a token bucket algorithm. The token bucket algorithm can use the user-configured values to determine the maximum rate of traffic allowed on an interface at a given moment in time. The token bucket algorithm is affected by all traffic entering or leaving (depending on where the traffic policy with Traffic Policing configured) and is useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream.

The token bucket algorithm provides users with three actions for each packet: a conform action, an exceed action, and an optional violate action. Traffic entering the interface with Traffic Policing configured is placed into one of these categories. Within these three categories, users can decide packet treatments. For instance, packets that conform can be configured to be transmitted, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

Traffic Policing is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In the most common Traffic Policing configurations, traffic that conforms is

transmitted and traffic that exceeds is sent with a decreased priority or is dropped. Users can change these configuration options to suit their network needs.

The Traffic Policing feature supports the following MIBs:

- CISCO-CLASS-BASED-QOS-MIB
- CISCO-CLASS-BASED-QOS-CAPABILITY-MIB

This feature also supports RFC 2697, *A Single Rate Three Color Marker*.

For information on how to configure the Traffic Policing feature, see the "Configuring Traffic Policing" module.

- Benefits of Traffic Policing, page 8
- Restrictions for Traffic Policing, page 8
- Prerequisites for Traffic Policing, page 9

Benefits of Traffic Policing

Bandwidth Management Through Rate Limiting

Traffic policing allows you to control the maximum rate of traffic sent or received on an interface. Traffic policing is often configured on interfaces at the edge of a network to limit traffic into or out of the network. Traffic that falls within the rate parameters is sent, whereas traffic that exceeds the parameters is dropped or sent with a different priority.

Packet Marking Through IP Precedence, QoS Group, and DSCP Value Setting

Packet marking allows you to partition your network into multiple priority levels or classes of service (CoS), as follows:

- Use traffic policing to set the IP precedence or differentiated services code point (DSCP) values for packets entering the network. Networking devices within your network can then use the adjusted IP Precedence values to determine how the traffic should be treated. For example, the DWRED feature uses the IP Precedence values to determine the probability that a packet will be dropped.
- Use traffic policing to assign packets to a QoS group. The router uses the QoS group to determine how to prioritize packets.

Restrictions for Traffic Policing

The following restrictions apply to the Traffic Policing feature:

- On a Cisco 7500 series router, traffic policing can monitor CEF switching paths only. In order to use the Traffic Policing feature, CEF must be configured on both the interface receiving the packet and the interface sending the packet.
- On a Cisco 7500 series router, traffic policing cannot be applied to packets that originated from or are destined to a router.
- Traffic policing can be configured on an interface or a subinterface.
- Traffic policing is not supported on the following interfaces:
  - Fast EtherChannel
  - Tunnel
  - PRI

- Any interface on a Cisco 7500 series router that does not support CEF

Prerequisites for Traffic Policing

On a Cisco 7500 series router, CEF must be configured on the interface before traffic policing can be used.

Traffic Shaping to Regulate Packet Flow

Regulating the packet flow (that is, the flow of traffic) on the network is also known as traffic shaping. Traffic shaping allows you to control the speed of traffic leaving an interface. This way, you can match the flow of the traffic to the speed of the interface receiving the packet.

Cisco provides three mechanisms for regulating or shaping traffic: Class-Based Traffic Shaping, Generic Traffic Shaping (GTS), and Frame Relay Traffic Shaping (FRTS).

For more information about traffic shaping, see the "Regulating Packet Flow Using Traffic Shaping" module.

For information on configuring Frame Relay and FRTS, see the "Configuring Frame Relay" module and the "MQC-Based Frame Relay Traffic Shaping" module, respectively.

Configuring Traffic Policing

This feature module describes the Traffic Policing feature. Traffic policing allows you to control the maximum rate of traffic transmitted or received on an interface.

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

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.

Feature Overview

The Traffic Policing feature performs the following functions:

- Limits the input or output transmission rate of a class of traffic based on user-defined criteria
- Marks packets by setting the ATM Cell Loss Priority (CLP) bit, Frame Relay Discard Eligibility (DE) bit, IP precedence value, IP differentiated services code point (DSCP) value, MPLS experimental value, and Quality of Service (QoS) group.

The Traffic Policing feature is applied when you attach a traffic policy contain the Traffic Policing configuration to an interface. A traffic policy is configured using the Modular Quality of Service (QoS) Command-Line Interface (CLI) (MQC).

The table below lists the feature history.

Table 1 Feature History

Cisco IOS Release	Enhancement
12.1(5)T	This command was introduced for Cisco IOS Release 12.1 T. A new Traffic Policing algorithm was introduced. The **violate-action** option became available. This feature became available on Cisco 2600, 3600, 4500, 7200, and 7500 series routers.
12.2(2)T	The **set-clp-transmit** option for the *action* argument was added to the **police** command. The **set-frde-transmit** option for the *action* argument was added to the **police** command. However, the **set-frde-transmit** option is not supported for Any Transport over Multiprotocol Label Switching (MPLS) (AToM) traffic in this release. The **set-mpls-exp-transmit** option for the *action* argument was added to the **police** command.
Cisco IOS	For information about feature support in Cisco IOS software, use Cisco Feature Navigator.

- Benefits, page 12
- Restrictions, page 13

Benefits

Bandwidth Management Through Rate Limiting

Traffic policing allows you to control the maximum rate of traffic transmitted or received on an interface. Traffic policing is often configured on interfaces at the edge of a network to limit traffic into or out of the network. In most Traffic Policing configurations, traffic that falls within the rate parameters is transmitted, whereas traffic that exceeds the parameters is dropped or transmitted with a different priority.

Packet Marking

Packet marking allows you to partition your network into multiple priority levels or classes of service (CoS). A packet is marked and these markings can be used to identify and classify traffic for downstream devices. In some cases, such as ATM Cell Loss Priority (CLP) marking or Frame Relay Discard Eligibility (DE) marking, the marking is used to classify traffic.

- Use traffic policing to set the IP precedence or DSCP values for packets entering the network. Networking devices within your network can then use the adjusted IP precedence values to determine how the traffic should be treated. For example, the Weighted Random Early Detection (WRED) feature uses the IP precedence values to determine the probability that a packet will be dropped.
- Use traffic policing to assign packets to a QoS group. The router uses the QoS group to determine how to prioritize packets within the router.

Traffic can be marked without using the Traffic Policing feature. If you want to mark traffic but do not want to use Traffic Policing, see the "Marking Network Traffic" module.

Packet Prioritization for Frame Relay Frames

The Traffic Policing feature allows users to mark the Frame Relay DE bit of the Frame Relay frame. The Frame Relay DE bit is one bit and, therefore, can be set to either 0 or 1. In congested environments, frames with the DE bit set to 1 are discarded before frames with the DE bit set to 0.

Packet Prioritization for ATM Cells

The Traffic Policing feature allows users to mark the ATM CLP bit in ATM cells. The ATM CLP bit is used to prioritize packets in ATM networks. The ATM CLP bit is one bit and, therefore, can be set to either 0 or 1. In congested environments, cells with the ATM CLP bit set to 1 are discarded before cells with the ATM CLP bit set to 0.

Restrictions

- On a Cisco 7500 series router, traffic policing can monitor Cisco Express Forwarding (CEF) switching paths only. In order to use the Traffic Policing feature, Cisco Express Forwarding must be configured on both the interface receiving the packet and the interface sending the packet.
- On a Cisco 7500 series router, traffic policing cannot be applied to packets that originated from or are destined to a router.
- Traffic policing can be configured on an interface or a subinterface.
- Traffic policing is not supported on the following interfaces:
  - Fast EtherChannel
  - Tunnel

Note Traffic policing is supported on tunnels that are using the Cisco generic routing encapsulation (GRE) tunneling protocol.

-
  - PRI
  - Any interface on a Cisco 7500 series router that does not support Cisco Express Forwarding

Supported Platforms

- Cisco 2500 series

Note Cisco IOS Release 12.2(2)T or later does not run on Cisco 2500 series routers.

- Cisco 2600 series
- Cisco 3640 routers
- Cisco 4500 series
- Cisco 7000 series with RSP7000
- Cisco 7100 series
- Cisco 7200 series
- Cisco 7500 series

Note To use the **set-clp-transmit** action available with this feature, the Enhanced ATM Port Adapter (PA-A3) is required. Therefore, the **set-clp-transmit**action is not supported on any platform that does not support the PA-A3 adapter (such as the Cisco 2600 series router, the Cisco 3640 router, and the 4500 series router). For more information, see the documentation for your specific router.

Supported Standards MIBs and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

Class-Based Quality of Service MIB

- CISCO-CLASS-BASED-QOS-MIB
- CISCO-CLASS-BASED-QOS-CAPABILITY-MIB

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

RFCs

- RFC 2697, *A Single Rate Three Color Marker*

Prerequisites

On a Cisco 7500 series router, Cisco Express Forwarding (CEF) must be configured on the interface before traffic policing can be used.

Configuration Tasks

- Configuring Traffic Policing, page 15
- Troubleshooting Tips, page 15

Configuring Traffic Policing

Command	Purpose
Router(config-pmap-c)# **police** *bps burst-normal burst-max* **conform-action** *action* **exceed-action** *action* **violate-action** *action*	Specifies a maximum bandwidth usage by a traffic class. **Note** The Traffic Policing feature works with a token bucket mechanism. There are currently two types of token bucket algorithms: a single token bucket algorithm and a two token bucket algorithm. A single token bucket system is used when the **violate-action** option is not specified, and a two token bucket system is used when the **violate-action** option is specified.

Troubleshooting Tips

- Check the interface type. Verify that your interface is not mentioned in the nonsupported interface description in the Restrictions, page 13 section of this module.
- For input traffic policing on a Cisco 7500 series router, verify that CEF is configured on the interface where traffic policing is configured.
- For output traffic policing on a Cisco 7500 series router, ensure that the incoming traffic is CEF-switched. Traffic policing cannot be used on the switching path unless CEF switching is enabled.

Monitoring and Maintaining Traffic Policing

Command	Purpose
Router# **show policy-map**	Displays all configured policy maps.
Router# **show policy-map** *policy-map-name*	Displays the user-specified policy map.
Router# **show policy-map interface**	Displays statistics and configurations of all input and output policies that are attached to an interface.

Configuration Examples

- Example Configuring a Service Policy that Includes Traffic Policing, page 16

Example Configuring a Service Policy that Includes Traffic Policing

The following configuration shows how to define a traffic class (with the **class-map**command) and associate that traffic class with a traffic policy (with the **policy-map** command). Traffic policing is applied in the traffic policy. The **service-policy** command is then used to attach the traffic policy to the interface.

In this example, traffic policing is configured with the average rate at 8000 bits per second, the normal burst size at 1500 bytes, and the excess burst size at 4000 bytes. Packets coming into Fast Ethernet interface 0/0 are evaluated by the token bucket algorithm to analyze whether packets conform, exceed, or violate the parameters. Packets that conform are transmitted, packets that exceed are assigned a QoS group value of 4 and are transmitted, packets that violate are dropped.

```
Router(config)# class-map acgroup2
Router(config-cmap)# match access-group 2
Router(config-cmap)# exit
Router(config)# policy-map police
Router(config-pmap)# class acgroup2
Router(config-pmap-c)# police 8000 1500 4000 conform-action transmit exceed-action set-
qos-transmit 4 violate-action drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface fastethernet 0/0
Router(config-if)# service-policy input police
```

MQC Traffic Shaping Overhead Accounting for ATM

The MQC Traffic Shaping Overhead Accounting for ATM feature enables a broadband aggregation system (BRAS) to account for various encapsulation types when applying quality of service (QoS) functionality to packets. Typically, in Ethernet digital subscriber line (DSL) environments, the encapsulation from the router to the digital subscriber line access multiplexer (DSLAM) is Gigabit Ethernet and the encapsulation from the DSLAM to the customer premises equipment (CPE) is ATM. ATM overhead accounting enables the router to account for ATM encapsulation on the subscriber line and for the overhead added by cell segmentation. This functionality enables the service provider to prevent overruns at the subscriber line and ensures that the router executes QoS features on the actual bandwidth used by ATM packets.

Finding Feature Information in This Module

Your Cisco IOS software release may not support all of the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To reach links to specific feature documentation in this module and to see a list of the releases in which each feature is supported, use the Feature Information for MQC Traffic Shaping Overhead Accounting for ATM, page 28.

Finding Support Information for Platforms and Cisco IOS and Catalyst OS Software Images

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

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.

Prerequisites for Traffic Shaping Overhead Accounting for ATM

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

Restrictions for Traffic Shaping Overhead Accounting for ATM

- The encapsulation type used within a policy map and between the parent policy map and the child policy map (in a hierarchical policy map structure) must be consistent.

Note This restriction applies to the Cisco 10000 series router only. This restriction does not apply to the Cisco 7600 series router.

- You must attach a policy map that is configured with ATM overhead accounting to only an Ethernet interface (or an IP session on an Ethernet interface).

Information About Traffic Shaping Overhead Accounting for ATM

Benefits of Traffic Shaping Overhead Accounting for ATM

The Traffic Shaping Overhead Accounting for ATM feature enables the broadband aggregation system (BRAS) to account for various encapsulation types when applying QoS to packets. Typically, in Ethernet digital subscriber line (DSL) environments, the encapsulation from the BRAS to the DSLAM is Gigabit Ethernet and the encapsulation from the DSLAM to the CPE is ATM. ATM overhead accounting enables the BRAS to account for ATM encapsulation on the subscriber line and for the overhead added by cell segmentation. This functionality enables the service provider to prevent overruns at the subscriber line and ensures that the router executes QoS features on the actual bandwidth used by ATM subscriber traffic.

BRAS and Encapsulation Types

Broadband aggregation system (BRAS) uses the encapsulation type that is configured for the DSLAM-CPE side to calculate the ATM overhead per packet.

DSLAM-CPE encapsulation types are based on Subnetwork Access Protocol (SNAP) and multiplexer (MUX) formats of ATM adaptation layer 5 (AAL5), followed by routed bridge (RBE), x-1483, x-dot1q-rbe, IP, PPP over Ethernet (PPPoE), or PPP over ATM (PPPoA) encapsulations. Because the DSLAM treats IP and PPPoE packets as payload, the BRAS does not account for IP and PPPoE encapsulations.

On the BRAS-DSLAM side, encapsulation is IEEE 802.1Q VLAN or Q-in-Q (qinq). However, because the DSLAM removes the BRAS-DSLAM encapsulation, the BRAS does not account for 802.1Q or qinq encapsulation.

AAL5 segmentation processing adds the additional overhead of the 5-byte cell headers, the AAL5 Common Part Convergence Sublayer (CPCS) padding, and the AAL5 trailer. For more information, see the ATM Overhead Calculation, page 19.

Subscriber Line Encapsulation Types

The router supports the following subscriber line encapsulation types:

- snap-rbe
- mux-rbe
- snap-dot1q-rbe
- mux-dot1q-rbe
- snap-pppoa
- mux-pppoa
- snap-1483routed
- mux-1483routed

Note The encapsulation types listed above are for AAL5, qinq, and dot1q encapsulations. User-defined encapsulations with offsets based on the platform in use are also supported. (For the Cisco 10000 series router, valid offsets are -63 to +63. For the Cisco 7600 series router, valid offsets are -48 to +48.)

ATM Overhead Calculation

The Traffic Shaping Overhead Accounting for ATM feature prevents oversubscription of a subscriber line by accounting for the ATM encapsulation overhead at the BRAS. When calculating the ATM overhead, the Traffic Shaping Overhead Accounting for ATM feature considers the following:

- The encapsulation type used by the BRAS
- The CPCS trailer overhead
- The encapsulation type used between the DSLAM and the CPE

The offset size (a parameter used to calculate ATM overhead accounting) is calculated using the following formula:

Offset size in bytes = (CPCS trailer overhead) + (DSLAM to CPE) - (BRAS encapsulation type)

This offset size, along with the packet size and packet assembler/disassembler (PAD) byte overhead in the CPCS, is used by the router to calculate the ATM overhead accounting rate.

Note A CPCS trailer overhead of 8 bytes corresponds to AAL5. A CPCS trailer overhead of 4 bytes corresponds to AAL3, but AAL3 is not supported.

Table 2 Offset Sizes, in Bytes, Used for ATM Overhead Calculation

Encapsulation Type in Use	BRAS	CPCS Trailer Overhead	DSLAM to CPE	Offset Size
dot1q mux-1483routed	18	8	3	-7
dot1q snap-1483routed	18	8	6	-4
dot1q mux-rbe	18	8	14	4
dot1q snap-rbe	18	8	24	14
dot1q mux-dot1q-rbe	18	8	18	8
dot1q snap-dot1q-rbe	18	8	28	18
qot1q mux-pppoa	18 + 6	8	2	-14
qot1q snap-pppoa	18 + 6	8	4	-12
qinq mux-1483routed	22	8	3	-11
qinq snap-1483routed	22	8	6	-8
qinq mux-rbe	22	8	14	0
qinq snap-rbe	22	8	24	10
qinq mux-dot1q-rbe	22	8	18	4
qing snap-dot1q-rbe	22	8	28	14
qinq mux-pppoa	22 + 6	8	2	-18
qinq snap-pppoa	22 + 6	8	4	-16

ATM Overhead Accounting and Hierarchical Policies

In hierarchical policies, you can enable ATM overhead accounting for shaping and bandwidth on parent policies and child policies. You are not required to enable ATM overhead accounting on a traffic class that does not contain the **bandwidth** or **shape** command. If you enable ATM overhead accounting on a child policy, then you must enable ATM overhead accounting on the parent policy. The parent and child classes must specify the same encapsulation type when ATM overhead accounting is enabled.

How to Configure Traffic Shaping Overhead Accounting for ATM

Configuring Traffic Shaping Overhead Accounting for ATM in a Hierarchical Policy

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **policy-map** *policy-map-name*
4. **class** *class-map-name*
5. **bandwidth** {*bandwidth-kbps* |**percent***percentage* | **remainingpercent***percentage*} [**account**{**qinq**| **dot1q**} [**aal5**] {*subscriber-encapsulation* | **user-defined***offset*}]
6. **bandwidth remaining ratio** *ratio* [**account** {**qinq** | **dot1q**} [**aal5**] {*subscriber-encapsulation* | **user-defined***offset*}]
7. **shape** [**average** |**peak**] *mean-rate* [*burst-size*] [*excess-burst-size*][**account**{**qinq**| **dot1q**} [**aal5**] {*subscriber-encapsulation* | **user-defined***offset*}]
8. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**policy-map** *policy-map-name* **Example:** `Router(config)# policy-map Business`	Creates or modifies the child policy and enters policy-map configuration mode. • Enter the policy map name. This is the name of the child policy and can be a maximum of 40 alphanumeric characters.
Step 4	**class** *class-map-name* **Example:** `Router(config-pmap)# class video`	Assigns the traffic class that you specify for the policy map and enters policy-map class configuration mode. • Enter the traffic class name. This is the name of the previously configured class map and can be a maximum of 40 alphanumeric characters.
Step 5	**bandwidth** {*bandwidth-kbps* | **percent***percentage* | **remainingpercent***percentage*} [**account**{**qinq**| **dot1q**} [**aal5**] {*subscriber-encapsulation* | **user-defined***offset*}] **Example:** `Router(config-pmap-c)# bandwidth 8000 account dot1q aal5 snap-pppoa`	Enables Class-Based Weighted Fair Queuing (CBWFQ) on the basis of the keywords and arguments specified, such as the following: • *bandwidth-kbps* --Specifies or modifies the minimum bandwidth allocated for a class that belongs to a policy map. Valid values are from 8 to 2488320, which represents from 1 to 99 percent of the link bandwidth. • **percent** *percentage* --Specifies or modifies the minimum percentage of the link bandwidth allocated for a class that belongs 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 that belongs to a policy map. Valid values are from 1 to 99. • **account** --Enables ATM overhead accounting. • **qinq** --Specifies queue-in-queue encapsulation as the BRAS-DSLAM encapsulation type. • **dot1q** --Specifies IEEE 802.1Q VLAN encapsulation as the BRAS-DSLAM encapsulation type. • **aal5** --Specifies the ATM adaptation layer 5 that supports connection-oriented variable bit rate (VBR) services. • *subscriber-encapsulation* --Specifies the encapsulation type at the subscriber line. For more information, see the Subscriber Line Encapsulation Types, page 19. • **user-defined** --Specifies the offset size that the router uses when calculating the ATM overhead. • *offset* --Specifies the offset size when calculating ATM overhead. Valid values are from -63 to +63 bytes. **Note** For the Cisco 7600 series router, valid values are from -48 to +48 bytes.

Command or Action	Purpose
Step 6 **bandwidth remaining ratio** *ratio* [**account** {**qinq** | **dot1q**} [**aal5**] {*subscriber-encapsulation* | **user-defined***offset*}] **Example:** `Router(config-pmap-c)# bandwidth remaining ratio 10 account dot1q aal5 snap-pppo`	(Optional) Specifies the bandwidth-remaining ratio for the subinterface along with ATM accounting parameters: • *ratio* --Specifies the bandwidth-remaining ratio for the subinterface. Valid values are 1 to 100. The default value is 1. **Note** For the Cisco 7600 series router, valid values are from 1 to 10000. The default value is 1. • **account** --Enables ATM overhead accounting. • **qinq** --Specifies queue-in-queue encapsulation as the BRAS-DSLAM encapsulation type. • **dot1q** --Specifies IEEE 802.1Q VLAN encapsulation as the BRAS-DSLAM encapsulation type. • **aal5** --Specifies the ATM adaptation layer 5 that supports connection-oriented VBR services. • *subscriber-encapsulation* --Specifies the encapsulation type at the subscriber line. For more information, see the Subscriber Line Encapsulation Types, page 19. • **user-defined** --Specifies the offset size that the router uses when calculating the ATM overhead. • *offset* --Specifies the offset size, in bytes, when calculating ATM overhead. Valid values are from -63 to +63. **Note** For the Cisco 7600 series router, valid values are from -48 to +48.

	Command or Action	Purpose
Step 7	**shape** [**average** |**peak**] *mean-rate* [*burst-size*] [*excess-burst-size*][**account**{**qinq**| **dot1q**} [**aal5**] {*subscriber-encapsulation* | **user-defined***offset*}] **Example:** `Router(config-pmap-c)# shape 8000 account qinq aal5 snap-dot1q-rbe`	Shapes traffic to the indicated bit rate and enables ATM overhead accounting on the basis of the keywords and arguments specified, such as the following: • **average** --(Optional) The committed burst (Bc) that specifies the maximum number of bits sent out in each interval. • **peak** --(Optional) Specifies the maximum number of bits sent out in each interval (the Bc + excess burst [Be]). The Cisco 10000 router and the SIP400 (on the Cisco 7600 series router) do not support this option. • *mean-rate* --Also called committed information rate (CIR). Indicates the bit rate used to shape the traffic, in bits per second. • *burst-size* --(Optional) The number of bits in a measurement interval (Bc). • *excess-burst-size* --(Optional) The acceptable number of bits permitted to go over the Be. • **account** --Enables ATM overhead accounting. • **qinq** --Specifies queue-in-queue encapsulation as the BRAS-DSLAM encapsulation type. • **dot1q** --Specifies IEEE 802.1Q VLAN encapsulation as the BRAS-DSLAM encapsulation type. • **aal5** --The ATM adaptation layer 5 that supports connection-oriented variable bit rate (VBR) services. • *subscriber-encapsulation* --Specifies the encapsulation type at the subscriber line. For more information, see the Subscriber Line Encapsulation Types, page 19. • **user-defined** --Specifies the offset size that the router uses when calculating the ATM overhead. • *offset* --Specifies the offset size when calculating ATM overhead. Valid values are from -63 to +63 bytes. **Note** For the Cisco 7600 series router, valid values are from -48 to +48 bytes.
Step 8	**end** **Example:** `Router(config-pmap-c)# end`	Exits policy-map class configuration mode.

Verifying the Configuration of Traffic Shaping Overhead Accounting for ATM

SUMMARY STEPS

1. **enable**
2. **show policy-map** [*policy-map-name*]
3. **show policy-map session**
4. **show running-config**
5. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**show policy-map** [*policy-map-name*] **Example:** `Router# show policy-map unit-test`	(Optional) Displays the configuration of all classes for a specified policy map or of all classes for all existing policy maps. • (Optional) Enter the policy map name. The name can be a maximum of 40 alphanumeric characters.
Step 3	**show policy-map session** **Example:** `Router# show policy-map session`	(Optional) Displays the QoS policy map in effect for a IPoE/PPPoE session.
Step 4	**show running-config** **Example:** `Router# show running-config`	(Optional) Displays the contents of the currently running configuration file.
Step 5	**exit** **Example:** `Router# exit`	Exits privileged EXEC mode.

Configuration Examples for Traffic Shaping Overhead Accounting for ATM

• Example Enabling Traffic Shaping Overhead Accounting for ATM, page 26

- Example Verifying Traffic Shaping Overhead Accounting for ATM, page 26

Example Enabling Traffic Shaping Overhead Accounting for ATM

In the following example, overhead accounting is enabled for bandwidth on the gaming and class-default class of the child policy map named subscriber_classes and on the class-default class of the parent policy map named subscriber_line. The voip and video classes do not have accounting explicitly enabled; these classes have ATM overhead accounting implicitly enabled because the parent policy has overhead accounting enabled. Notice that the features in the parent and child policies use the same encapsulation type.

```
policy-map subscriber_classes
 class voip
  priority level 1
  police 8000
 class video
  priority level 2
  police 8000
 class gaming
  bandwidth remaining percent 80 accountaal5 snap-rbe-dot1q
 class class-default
  bandwidth remaining percent 20 accountaal5 snap-rbe-dot1q
policy-map subscriber_line
 class class-default
  bandwidth remaining ratio 10 accountaal5 snap-rbe-dot1q
  shape average 512 account aal5snap-rbe-dot1q
  service policy subscriber_classes
```

Example Verifying Traffic Shaping Overhead Accounting for ATM

```
Router# show policy-map interface
```

```
Router# show policy-map session output

SSS session identifier 2 -
Service-policy output:  ATM_OH_POLICY
    Class-map: class-default (match-any)
      0 packets, 0 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any
      Queueing
      queue limit 2500 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0
      shape (average) cir 10000000, bc 40000, be 40000
      target shape rate 10000000
       Overhead Accounting Enabled
```

The following output from the **show running-config** command indicates that ATM overhead accounting is enabled for shaping. The BRAS-DSLAM encapsulation is dot1q and the subscriber line encapsulation is snap-rbe based on the AAL5 service.

```
subscriber policy recording rules limit 64
no mpls traffic-eng auto-bw timers frequency 0
call rsvp-sync
!
controller T1 2/0
framing sf
linecode ami
!
controller T1 2/1
framing sf
linecode ami
```

```
!
!
policy-map unit-test
class class-default
shape average percent 10 account dot1q aal5 snap-rbe
!
```

Additional References

The following sections provide references related to traffic shaping overhead accounting for ATM.

Related Documents

Related Topic	Document Title
QoS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	*Cisco IOS Quality of Service Solutions Command Reference*
Modular Quality of Service (QoS) Command-Line Interface (CLI) (MQC), hierarchical policies, policy maps	"Applying QoS Features Using the MQC" module
Policing and shaping traffic	"Policing and Shaping Overview" module

Standards

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

MIBs

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

RFCs

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

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Command Reference

The following commands are introduced or modified in the feature or features documented in this module. For information about these commands, see the Cisco IOS Quality of Service Solutions Command Reference. For information about all Cisco IOS commands, use the Command Lookup Tool at http://tools.cisco.com/Support/CLILookup or a Cisco IOS master commands list.

- **bandwidth (policy-map class)**
- **bandwidth remaining ratio**
- **shape (policy-map class)**
- **show policy-map interface**
- **show policy-map session**
- **show running-config**

Feature Information for MQC Traffic Shaping Overhead Accounting for ATM

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

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.

Table 3 Feature Information for MQC Traffic Shaping Overhead Accounting for ATM

Feature Name	Releases	Feature Information
MQC Traffic Shaping Overhead Accounting for ATM	12.2(31)SB2 12.2(33)SRC 12.2(33)SB	The MQC Traffic Shaping Overhead Accounting for ATM feature enables a broadband aggregation system (BRAS) to account for various encapsulation types when applying QoS functionality to packets. In Release 12.2(31)SB2, this feature was introduced and implemented on the Cisco 10000 series router for the PRE3. In Release 12.2(33)SRC, support was added for the Cisco 7600 series router. In Release 12.2(33)SB, support was added for the Cisco 7300 series router. The following commands were introduced or modified: **bandwidth (policy-map class)**, **bandwidth remaining ratio**, **shape (policy-map class)**, **show policy-map interface**, **show policy-map session**, **show running-config**.

Two-Rate Policer

This document describes the Two-Rate Policer feature and how to configure it. Two-Rate Policer allows you to manage traffic rates through an interface; it is especially helpful in managing network bandwidth where large packets are in the same traffic stream.

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.

Prerequisites for Two-Rate Policer

Supported Platforms

- Cisco 2600 series
- Cisco 3620
- Cisco 3640
- Cisco 7100 series
- Cisco 7200 series
- Cisco 7500 series (VIP-based platform only)

Note The **set-clp-transmit** action available with Two-Rate Policer, the Enhanced ATM Port Adapter (PA-A3) is required. The **set-clp-transmit** action is not supported on any platform that does not support the PA-A3 adapter (such as the Cisco 2600 series router, the Cisco 3620 router, and the 3640 router). For more information, see the documentation for your specific router.

- On a Cisco 7500 series router, Cisco Express Forwarding or Distributed Cisco Express Forwarding must be configured on the interface before you can use the Two-Rate Policer.
- A traffic class and a service policy must be created, and the service policy must be attached to a specified interface. These tasks are performed using the Modular quality of service (QoS) Command-Line Interface (CLI) (MQC). For information on the MQC, see the "Applying QoS Features Using the MQC" module.

Restrictions for Two-Rate Policer

The following restrictions apply to the Two-Rate Policer feature:

- On a Cisco 7500 series router, traffic policing can monitor Cisco Express Forwarding or Distributed Cisco Express Forwarding switching paths only. Cisco Express Forwarding or Distributed Cisco Express Forwarding must be configured on both the interface receiving the packet and the interface sending the packet.
- On a Cisco 7500 series router, traffic policing cannot be applied to packets that originated from or are destined to a router.
- Two-rate policing can be configured on an interface, a subinterface, a Frame Relay data-link connection identifier (DLCI), and an ATM permanent virtual circuit (PVC).
- Two-rate policing is not supported on the following interfaces:
  - Fast EtherChannel
  - PRI
  - Any interface on a Cisco 7500 series router that does not support Cisco Express Forwarding or Distributed Cisco Express Forwarding

Information About Two-Rate Policer

Networks police traffic by limiting the input or output transmission rate of a class of traffic based on user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or class of service (CoS).

The Two-Rate Policer performs the following functions:

- Limits the input or output transmission rate of a class of traffic based on user-defined criteria.
- Marks packets by setting the IP precedence value, IP differentiated services code point (DSCP) value, Multiprotocol Label Switching (MPLS) experimental value, Quality of Service (QoS) group, ATM Cell Loss Priority (CLP) bit, and the Frame Relay Discard Eligibility (DE) bit.

With the Two-Rate Policer, you can enforce traffic policing according to two separate rates--committed information rate (CIR) and peak information rate (PIR). You can specify the use of these two rates, along with their corresponding values, by using two keywords, **cir** and **pir**, of the **police** command.

The Two-Rate Policer manages the maximum rate of traffic through a token bucket algorithm. The token bucket algorithm can use the user-configured values to determine the maximum rate of traffic allowed on

an interface at a given moment in time. The token bucket algorithm is affected by all traffic entering or leaving the interface (depending on the location of the interface on which the Two-Rate Policer is configured) and is useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream.

The token bucket algorithm provides users with three actions for each packet: a conform action, an exceed action, and an optional violate action. Traffic coming into the interface with the Two-Rate Policer configured is assigned one of these categories. Within these three categories, users can decide packet treatments. For instance, packets that conform can be configured to be sent, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

The Two-Rate Policer is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In the most common configurations, traffic that conforms is sent and traffic that exceeds is sent with a decreased priority or is dropped. Users can change these configuration options to suit their network needs.

Note Two-Rate Policer enables you to use Differentiated Services (DiffServ) Assured Forwarding (AF) Per-Hop Behavior (PHB) traffic conditioning. For more information about DiffServ, see the "Implementing DiffServ for End-to-End Quality of Service Overview" module.

Note Starting with Cisco IOS Release 12.1(5)T, you can police traffic by using the Traffic Policing feature (sometimes referred to as the single-rate policer). The Two-Rate Policer (available with Cisco IOS Release 12.2(4)T) is in addition to the Traffic Policing feature, and it provides additional functionality. For more information about the Traffic Policing feature, see the "Traffic Policing" module.

- Benefits, page 33

Benefits

Bandwidth Management Through Rate Limiting

Two-Rate Policer provides improved bandwidth management through rate limiting. Before this feature was available, you could police traffic with the single-rate Traffic Policing feature. The Traffic Policing feature provided a certain amount of bandwidth management by allowing you to set the peak burst size (be). The Two-Rate Policer supports a higher level of bandwidth management and supports a sustained excess rate. With the Two-Rate Policer, you can enforce traffic policing according to two separate rates--CIR and PIR--specified in bits per second (bps).

Packet Marking Through IP Precedence, DSCP Value, MPLS Experimental Value, and the QoS Group Setting

In addition to rate-limiting, the Two-Rate Policer allows you to independently mark the packet according to whether the packet conforms, exceeds, or violates a specified rate. Packet marking also allows you to partition your network into multiple priority levels or CoSs.

- Use the Two-Rate Policer to set the IP precedence value, the IP DSCP value, or the MPLS experimental value for packets that enter the network. Then networking devices within your network can use this setting to determine how the traffic should be treated. For example, the Weighted Random

Early Detection (WRED) feature uses the IP precedence value to determine the probability that a packet will be dropped.

- Use the Two-Rate Policer to assign packets to a QoS group. The router uses the QoS group to determine how to prioritize packets within the router.

If you want to mark traffic but do not want to use the Two-Rate Policer, see the "Marking Network Traffic" module.

Packet Marking for Frame Relay Frames

The Two-Rate Policer allows users to mark the Frame Relay DE bit of the Frame Relay frame. The Frame Relay DE bit is one bit and, therefore, can be set to either 0 or 1. In congested environments, frames that have the DE bit set to 1 are discarded before frames that have the DE bit set to 0.

Packet Marking for ATM Cells

The Two-Rate Policer allows users to mark the ATM CLP bit in ATM cells. The ATM CLP bit is used to prioritize packets in ATM networks. The ATM CLP bit is one bit and, therefore, can be set to either 0 or 1. In congested environments, cells that have the ATM CLP bit set to 1 are discarded before cells that have the ATM CLP bit set to 0.

How to Use the Two-Rate Policer

Configuring the Two-Rate Policer

Command	Purpose
`Router(config-pmap-c)#` **police cir** *cir* [**bc** *conform-burst*] **pir** *pir* [**be** *peak-burst*	Specifies that both the CIR and the PIR are to be used for two-rate traffic policing. The **bc** and **be** keywords and their associated arguments (*conform-burst* and *peak-burst*, respectively) are optional. Specifies the action taken on a packet when you enable an optional **action** argument. **Note** The Two-Rate Policer works by using a token bucket mechanism. There are currently two types of token bucket algorithms: a single token bucket algorithm (available through the Traffic Policing feature) and a two token bucket algorithm (available through the Two-Rate Policer).

Verifying the Two-Rate Policer Configuration

Command	Purpose
`Router#` **show policy-map interface**	Displays statistics and configurations of all input and output policies attached to an interface.

Troubleshooting Tips

- Check the interface type. Verify that your interface is not listed as a nonsupported interface in the Restrictions for Two-Rate Policer, page 32 section of this module.
- For input traffic policing on a Cisco 7500 series router, verify that Cisco Express Forwarding or Distributed Cisco Express Forwarding is configured on the interface on which traffic policing is configured.
- For output traffic policing on a Cisco 7500 series router, ensure that the incoming traffic is Cisco Express Forwarding-switched or Distributed Cisco Express Forwarding-switched. Traffic policing cannot be used on the switching path unless Cisco Express Forwarding or Distributed Cisco Express Forwarding switching is enabled.

Monitoring and Maintaining the Two-Rate Policer

Command	Purpose
`Router#` **show policy-map**	Displays all configured policy maps.
`Router#` **show policy-map** *policy-map-name*	Displays the user-specified policy map.
`Router#` **show policy-map interface**	Displays statistics and configurations of all input and output policies that are attached to an interface.

Configuration Examples

Example Limiting the Traffic Using a Policer Class

In this example, the Two-Rate Policer is configured on a class to limit traffic to an average committed rate of 500 kbps and a peak rate of 1 Mbps:

```
Router(config)# class-map
 police
Router(config-cmap)# match
 access-group 10
1
```

```
Router(config-cmap)# policy-map
 policy1
Router(config-pmap)# class
 police
Router(config-pmap-c)# police
 cir 500000 bc 10000 pir 1000000 be 10000 conform-action transmit exceed-action set-prec-
transmit 2 violate-action drop
Router(config)# interface
 serial3/0
Router(config-if)# service-policy
 output policy1
Router(config-if)# end
Router# show
 policy-map policy1
 Policy Map policy1
  Class police
   police cir 500000 conform-burst 10000 pir 1000000 peak-burst 10000 conform-action
transmit exceed-action set-prec-transmit 2 violate-action drop
```

Traffic marked as conforming to the average committed rate (500 kbps) will be sent as is. Traffic marked as exceeding 500 kbps, but not exceeding 1 Mbps, will be marked with IP Precedence 2 and then sent. All traffic exceeding 1 Mbps will be dropped. The burst parameters are set to 10,000 bytes.

```
Router# show
 policy-map interface serial3/0
 Serial3/0
  Service-policy output: policy1
   Class-map: police (match all)
    148803 packets, 36605538 bytes
    30 second offered rate 1249000 bps, drop rate 249000 bps
    Match: access-group 101
    police:
     cir 500000 bps, conform-burst 10000, pir 1000000, peak-burst 100000
     conformed 59538 packets, 14646348 bytes; action: transmit
     exceeded 59538 packets, 14646348 bytes; action: set-prec-transmit 2
     violated 29731 packets, 7313826 bytes; action: drop
     conformed 499000 bps, exceed 500000 bps violate 249000 bps
   Class-map: class-default (match-any)
    19 packets, 1990 bytes
    30 seconds offered rate 0 bps, drop rate 0 bps
    Match: any
```

Additional References

The following sections provide references related to the Two-Rate Policer feature.

Related Documents

Related Topic	Document Title
MQC	• "Applying QoS Features Using the MQC" module
QoS features such as class-based weighted fair queueing (CBWFQ), traffic marking, and traffic policing	• "Configuring Weighted Fair Queueing" module • "Marking Network Traffic" module • "Traffic Policing" module

Standards

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

MIBs

MIB	MIBs Link
• CISCO-CLASS-BASED-QOS-MIB • CISCO-CLASS-BASED-QOS-CAPABILITY-MIB	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 2698	*A Two Rate Three Color Marker*

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Two-Rate Policer

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

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.

Table 4 Feature Information for Two-Rate Policer

Feature Name	Releases	Feature Information
Two-Rate Policer	12.2(4)T	This feature was introduced.
	12.2(4)T3	Support for the Cisco 7500 series routers was added.
	12.0(26)S	This feature was integrated into Cisco IOS Release 12.0(26)S for the Cisco 7200 and 7500 series routers.
	12.2(28)SB	This feature was integrated into Cisco IOS Release 12.2(28)SB.
	12.2(33)SRA	This feature was integrated into Cisco IOS Release 12.2(33)SRA.
	12.2(33)SXH	This feature was integrated into Cisco IOS Release 12.2(33)SXH.
	Cisco IOS XE 3.1.0 SG	This feature was integrated into Cisco IOS XE 3.1.0 SG.

Control Plane Policing

The Control Plane Policing feature allows you to configure a quality of service (QoS) filter that manages the traffic flow of control plane packets to protect the control plane of Cisco IOS routers and switches against reconnaissance and denial-of-service (DoS) attacks. In this way, the control plane (CP) can help maintain packet forwarding and protocol states despite an attack or heavy traffic load on the router or switch.

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.

Prerequisites for Control Plane Policing

The Modular Quality of Service (QoS) Command-Line interface (CLI) (MQC) is used to configure the packet classification and policing functionality of the Control Plane Policing feature.

Before configuring Control Plane Policing (CoPP), you should understand the procedures for using the MQC. For information about the MQC, see the "Applying QoS Features Using the MQC" module.

Restrictions for Control Plane Policing

Aggregate and Distributed Control Plane Policing

Aggregate policing is supported in Cisco IOS Release 12.0(29)S, Cisco IOS Release 12.2(18)S, Cisco IOS Release 12.3(4)T, and later releases.

Distributed policing is supported only in Cisco IOS Release 12.0(30)S and later Cisco IOS 12.0S releases.

Output Rate-Limiting Support

Output rate-limiting is performed in silent (packet discard) mode. Silent mode enables a router to silently discard packets using policy maps applied to output control plane traffic with the **service-policy output** command. For more information, see the Output Rate-Limiting and Silent Mode Operation, page 47.

Output rate-limiting (policing) in silent mode is supported only in:

- Cisco IOS Release 12.2(25)S and later Cisco IOS 12.2S releases
- Cisco IOS Release 12.3(4)T and later Cisco IOS 12.3T releases

Output rate-limiting is not supported for distributed control plane services in Cisco IOS 12.0S releases or in Cisco IOS 12.2SX releases.

Output rate-limiting is not supported on the Cisco 7500 series and Cisco 10720 Internet router.

MQC Restrictions

The Control Plane Policing feature requires the MQC to configure packet classification and policing. All restrictions that apply when you use the MQC to configure policing also apply when you configure control plane policing. Only two MQC actions are supported in policy maps--**police** and **drop**.

Note On the Cisco 10720 Internet router, only the **police**command, not the **drop** command, is supported in policy maps. In addition, in a QoS service policy that is attached to the Cisco 10720 control plane, the **police**command does not support **set** actions as arguments in **conform-action**, **exceed-action**, and **violate-action** parameters.

Features that require Network-Based Application Recognition (NBAR) classification may not work well at the control plane level. The following classification (match) criteria are supported on all platforms:

- Standard and extended IP access lists (ACLs).
- In class-map configuration mode: **match ip dscp**, **match ip precedence**, and **match protocol arp,**and **match protocol pppoe**commands.

Note In the Cisco IOS 12.2SX release, the **match protocol arp** command is not supported.

On the Cisco 10720 Internet router, the following MQC commands are also supported in class-map configuration mode: **match input-interface**, **match mpls experimental**, **match protocol ipv6**, and **match qos-group**. When using these commands for control plane policing on the Cisco 10720 Internet router, note the following restrictions:

- Packet classification using match criteria is not supported for packets that cannot be classified in the Cisco 10720 data path, such as unknown Layer 2 encapsulation and IP options.
- The following IPv6 fields are not be supported in packet classification for IPv6 QoS on the Cisco 10720 Internet router and are, therefore, not supported for control plane policing:
    - IPv6 source and destination addresses

- Layer 2 class of service (CoS)
- IPv6 routing header flag
- IPv6 undetermined transport flag
- IPv6 flow label
- IP Real-Time transport Protocol (RTP)

Note Packets that are not supported for QoS packet classification on the Cisco 10720 Internet router are not policed in the default traffic class for control plane policing.

CISCO-CLASS-BASED-QOS-MIB Control Plane Support

In Cisco IOS Release 12.3(7)T and later Cisco IOS 12.3T releases, the CISCO-CLASS-BASED-QOS-MIB is extended to manage control plane QoS policies and provide information about the control plane.

Cisco IOS Release 12.2(18)SXD1

In Cisco IOS Release 12.2(18)SXD1 and later releases, Hardware Control Plane Interface for Control Plane Policing has the following restrictions:

- Supported only with Supervisor Engine 720. Not supported with Supervisor Engine 2.
- Does not support CoPP output rate-limiting (policing).
- Does not support the CoPP silent operation mode.
- Cisco IOS Release 12.2(18)SXD1 and later releases automatically install the CoPP service policy on all DFC-equipped switching modules.

For more information about control plane policing in Cisco IOS Release 12.2(18)SXD1 and later releases, see either of these publications:

- For Catalyst 6500 series switches, see the "Configuring Control Plane Policing (CoPP)" module.
- For Cisco 7600 series routers, see the "Configuring Denial of Service Protection" module.

Information About Control Plane Policing

Benefits of Control Plane Policing

Configuring the Control Plane Policing feature on your Cisco router or switch provides the following benefits:

- Protection against DoS attacks at infrastructure routers and switches
- QoS control for packets that are destined to the control plane of Cisco routers or switches

- Ease of configuration for control plane policies
- Better platform reliability and availability

Terms to Understand

Because different platforms can have different architectures, the following set of terms is defined. The figure below illustrates how control plane policing works.

Figure 1 Layout of Control Plane, Central Switch Engine, Distributed Switch Engines, and Line Cards on a Router

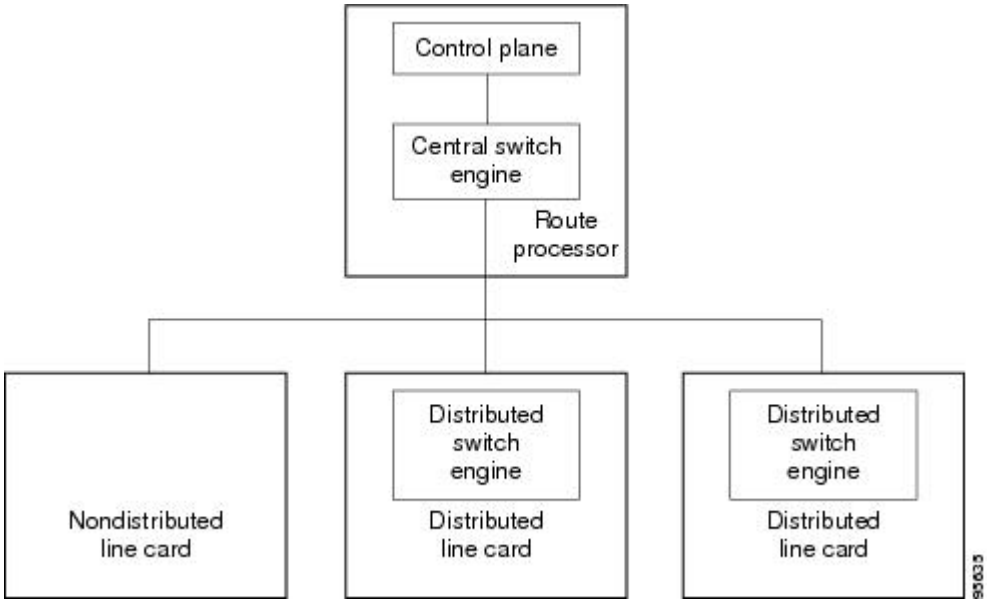

- Control plane (CP)--A collection of processes that run at the process level on the route processor (RP). These processes collectively provide high-level control for most Cisco IOS functions.
- Central switch engine--A device that is responsible for high-speed routing of IP packets. It also typically performs high-speed input and output services for nondistributed interfaces. (See nondistributed line cards.) The central switch engine is used to implement aggregate CP protection for all interfaces on the router.

Note All IP packets that are destined for the CP should pass through the central switch engine before they are forwarded to the process level.

On the Cisco 10720 Internet router, control plane policing is implemented on Cisco Parallel eXpress Forwarding (PXF) in a Toaster-based architecture. PXF is a hardware-based central switch engine that can filter traffic at a higher rate than the route processor. PXF switches all data traffic separately from the route processor. PXF packet processing occurs at an intermediate step between the nondistributed line cards and the route processor shown in the figure above. In addition to the regular punting, PXF also punts certain types of packets (such as unknown Layer 2 encapsulation and packets with IP options) to the RP for further processing at interrupt level.

Note On the Cisco 10720 Internet router, you can configure enhanced RP protection by using the **ip option drop** command to drop IPv4 packets with IP options that are punted to the RP by PXF. Tunneled IPv4 packets and IPv4 packets with an unsupported encapsulation method are not dropped. For more information, see the "ACL IP Options Selective Drop" module.

- Distributed switch engine--A device that is responsible for high-speed switching of IP packets on distributed line cards without using resources from the central switch engine. It also typically performs input and output services for the line card. Each distributed switch engine is used to implement distributed CP services for all ports on a line card. Input CP services distribute the processing load across multiple line cards and conserve vital central switch engine resources. Distributed CP services are optional; however, they provide a more refined level of service than aggregate services.
- Nondistributed line cards--Line cards that are responsible for receiving packets and occasionally performing input and output services. All packets must be forwarded to the central switch engine for a routing or switching decision. Aggregate CP services provide coverage for nondistributed line cards.

Note Distributed CP services are supported only in Cisco IOS Release 12.0(30)S and later 12.0S releases.

Control Plane Security and Packet QoS Overview

To protect the CP on a router from DoS attacks and to provide packet QoS, the Control Plane Policing feature treats the CP as a separate entity with its own ingress (input) and egress (output) ports, which are like ports on a router and switch. Because the Control Plane Policing feature treats the CP as a separate entity, a set of rules can be established and associated with the ingress and egress ports of the CP.

These rules are applied only after the packet has been determined to have the CP as its destination or when a packet exits the CP. Thereafter, you can configure a service policy to prevent unwanted packets from progressing after a specified rate limit has been reached; for example, a system administrator can limit all TCP/SYN packets that are destined for the CP to a maximum rate of 1 megabit per second.

Input CP services are executed after router input port services have been performed and after a routing decision on the input path has been made. As shown in the figure below, CP security and packet QoS are applied on:

Figure 2 Input Control Plane Services: Aggregate and Distributed Services

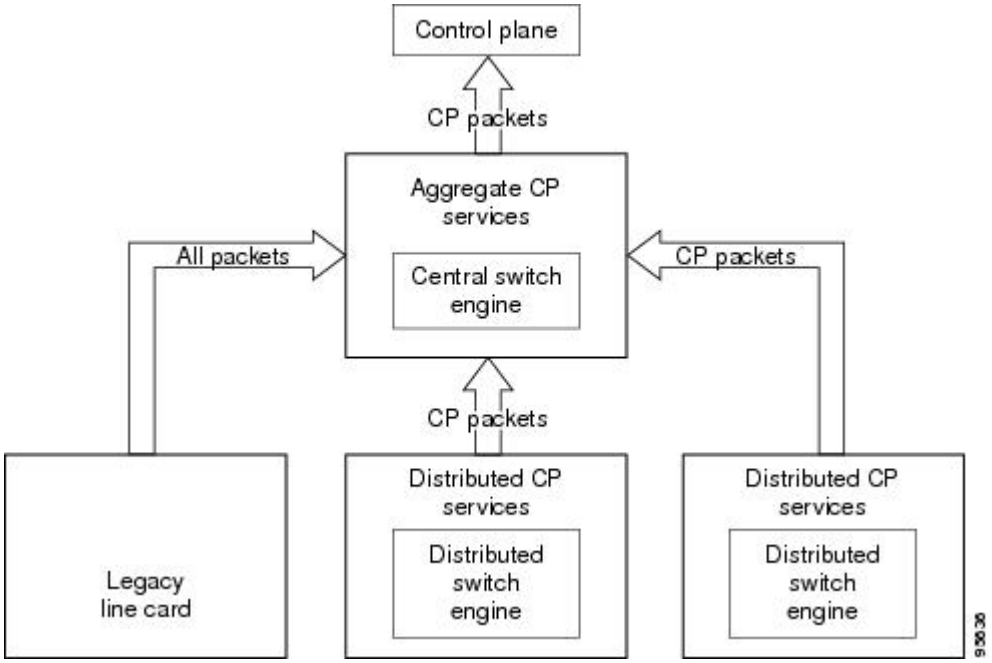

The following types of Layer 3 packets are forwarded to the control plane and processed by aggregate and distributed control plane policing:

- Routing protocol control packets
- Packets destined for the local IP address of the router
- Packets from management protocols (such as Simple Network Management Protocol [SNMP], Telnet, and secure shell [SSH])

Note Ensure that Layer 3 control packets have priority over other packet types that are destined for the control plane.

Aggregate Control Plane Services

Aggregate control plane services provide control plane policing for all CP packets that are received from all line-card interfaces on the router.

The central switch engine executes normal input port services and makes routing decisions for an incoming packet: if the packet is destined for the CP, aggregate services are performed. Because CP traffic from all line cards must pass through aggregate CP services, these services manage the cumulative amount of CP traffic that reaches the CP.

Aggregate CP service steps are as follows:

1 The line card receives a packet and delivers it to the central switch engine.

Note Before the packet is sent to the central switch engine, additional processing may be necessary for platforms that support hardware-level policing or platform-specific aggregate policing. It is possible that the packet may undergo multiple checks before it undergoes the generic Cisco IOS check.

1 The interfaces perform normal (interface-level) input port services and QoS.
2 The central switch engine performs Layer 3 switching or makes a routing decision, determining whether or not the packet is destined for the CP.
3 The central switch engine performs aggregate CP services for all CP packets.
4 On the basis of the results of the aggregate CP services, the central switch engine either drops the packet or delivers the packet to the CP for final processing.

Functionality Highlights of Aggregate CP Services

The following list highlights the functionality of aggregate CP services:

- Aggregate CP services are defined for a single input interface, such as the CP, and represent an aggregate for all ports on a router.
- Modular QoS is used to define CP services. Class maps and policy maps for both DoS protection and packet QoS are defined for a single aggregate CP service policy.
- Modular QoS does not prevent a single bad port from consuming all allocated bandwidth. Class maps that match an interface or subinterface may be able to constrain the contribution of each interface through an interface-specific policy map.

Distributed Control Plane Services

Distributed control plane services provide control plane policing for all CP packets that are received from the interfaces on a line card.

A distributed switch engine executes normal input port services and makes routing decisions for a packet: if the packet is destined for the CP, distributed CP services are performed. Afterwards, CP traffic from each line card is forwarded to the central switch engine where aggregate CP services are applied.

Note Distributed CP services may also forward conditioned packets to the central switch engine. In this case, aggregate CP services are also performed on the conditioned CP traffic.

Distributed CP service steps are as follows:

1 A line card receives a packet and delivers it to the distributed switch engine.
2 The distributed switch engine performs normal (interface-level) input port services and QoS.
3 The distributed switch engine performs Layer 2 or Layer 3 switching or makes a routing decision, determining whether the packet is destined for the CP.
4 The distributed switch engine performs distributed CP services for all CP packets.
5 On the basis of the results of the distributed CP services, the distributed switch engine either drops the packet or marks the packet and delivers it to the central switch engine for further processing.
6 The central switch engine performs aggregate CP services and delivers the packet to the CP for final processing.

Functionality Highlights of Distributed CP Services

The following list highlights the functionality of distributed CP services:

- Distributed CP services are defined for a single input interface, such as the distributed CP, and represent an aggregate for all ports on a line card.
- The MQC is used to define CP services. Class maps and policy maps for both DoS protection and packet QoS are defined for a single distributed CP service policy. Each line card may have a unique CP service policy that applies traffic classifications, QoS policies, and DoS services to packets received from all ports on the line card in an aggregate way.
- The MQC does not prevent one bad port from consuming all allocated bandwidth on a line card. Class maps that match an interface or subinterface may be able to constrain the contribution of each interface through an interface-specific policy map.
- Distributed CP services allow you to limit the number of CP packets forwarded from a line card to the central switch engine. The total number of CP packets received from all line cards on a router may exceed aggregate CP levels.

Usage of Distributed CP Services

The purpose of CP protection and packet QoS is to apply sufficient control to the packets that reach the control plane. To successfully configure this level of CP protection, you must:

- Apply traditional QoS services using the MQC to CP packets.
- Protect the path to the control plane against indiscriminate packet dropping due to resource exhaustion. If packets are not dropped according to user-defined QoS policies, but are dropped due to a resource limitation, the QoS policy is not maintained.

Distributed CP services allow you to configure specific CP services that are enforced at the line-card level and are required for the following reasons:

- While under a DoS attack, line-card resources may be consumed. In this case, you must configure a drop policy to identify important packets. The drop policy ensures that all important packets arrive to the central switch engine for aggregate CP protection and arrive later to the CP. Distributed CP services allow routers to apply the appropriate drop policy when resources are consumed and therefore maintain the desired QoS priorities. If a line card indiscriminately drops packets, the aggregate CP filter becomes ineffective and the QoS priorities are no longer maintained.
- It is not possible to prevent one interface from consuming all aggregate CP resources. A DoS attack on one port may negatively impact CP processing of traffic from other ports. Distributed CP services allow you to limit the amount of important traffic that is forwarded by a line card to the CP. For example, you can configure a layered approach in which the combined rates of all line cards are over-subscribed compared to the aggregate rate. The rate of each individual line card would be below the aggregate rate, but combined together, the rates of all line cards exceed it. This over-subscription model is commonly used for other resource-related functions and helps limit the contribution of CP packets from any one line card.
- Distributed CP services provide for slot-level (line-card) filtering. Customer-facing interfaces may have greater security requirements (with more restrictions or for billing reasons) than network-facing interfaces to backbone devices.
- Because distributed CP protection allows you to configure packet filters on a per-line-card basis, processing cycles on line cards may offload aggregate level processing. You can configure Border Gateway Protocol (BGP) filtering at the distributed level for interfaces that use BGP, allowing the aggregate level to filter packets with the remaining filter requirements. Or you can configure identical filters for distributed and aggregate CP services with a distributed packet marking scheme that informs

the aggregate filter that a packet has already been checked. Distributed CP service processing further reduces aggregate processing and can significantly reduce the load on aggregate CP services.

Output Rate-Limiting and Silent Mode Operation

A router is automatically enabled to silently discard packets when you configure output policing on control plane traffic using the **service-policy output** *policy-map-name* command.

Rate-limiting (policing) of output traffic from the CP is performed in silent mode. In silent mode, a router that is running Cisco IOS software operates without sending any system messages. If a packet that is exiting the control plane is discarded for output policing, you do not receive an error message.

When control plane policing is configured for output traffic, error messages are not generated in the following cases:

- Traffic that is being transmitted to a port to which the router is not listening
- A connection to a legitimate address and port that is rejected because of a malformed request

The silent mode functionality and output policing on CP traffic are supported only in:

- Cisco IOS Release 12.2(25)S and later Cisco IOS 12.2S releases
- Cisco IOS Release 12.3(4)T and later Cisco IOS 12.3T releases

Silent mode and output policing on CP traffic are not supported for distributed control plane services.

How to Use Control Plane Policing

Defining Aggregate Control Plane Services

To configure aggregate CP services, such as packet rate control and silent packet discard, for the active route processor, complete the following steps.

Before you enter control-plane configuration mode to attach an existing QoS policy to the control plane, you must first create the policy using MQC to define a class map and policy map for control plane traffic.

Note

- Platform-specific restrictions, if any, are checked when the service policy is applied to the control plane interface.
- Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases. (Note that output policing does not provide any performance benefits. It simply controls the information that is leaving the device.)

>

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **control-plane**
4. **service-policy** {**input**| **output** *policy-map-name*
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**control-plane** **Example:** `Router(config)# control-plane`	Enters control-plane configuration mode (a prerequisite for Defining Aggregate Control Plane Services, page 47).
Step 4	**service-policy** {**input**| **output** *policy-map-name* **Example:** `Router(config-cp)# service-policy input control-plane-policy`	Attaches a QoS service policy to the control plane. Note the following points: • **input** --Applies the specified service policy to packets received on the control plane. • **output** --Applies the specified service policy to packets transmitted from the control plane and enables the router to silently discard packets. • *policy-map-name* --Name of a service policy map (created using the **policy-map** command) to be attached. The name can be a maximum of 40 alphanumeric characters.
Step 5	**end** **Example:** `Router(config-cp)# end`	(Optional) Returns to privileged EXEC mode.

Defining Distributed Control Plane Services

To configure distributed CP services, such as packet rate control, for packets that are destined for the CP and sent from the interfaces on a line card, complete the following steps.

Before you enter control-plane configuration mode to attach an existing QoS policy for performing distributed control-plane services, you must first create the policy using MQC to define a class map and policy map for control-plane traffic.

Note

- Platform-specific restrictions, if any, are checked when the service policy is applied to the control plane interface.
- Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases. (Note that output policing does not provide any performance benefits. It simply controls the information that is leaving the device.)
- With Cisco IOS 12.2SX releases, the Supervisor Engine 720 automatically installs the service policy on all DFC-equipped switching modules.

>

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **control-plane** [**slot** *slot-number*]
4. **service-policy input** *policy-map-name*
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**control-plane** [**slot** *slot-number*] **Example:** `Router(config)# control-plane slot 3`	Enters control-plane configuration mode, which allows you to optionally attach a QoS policy (used to manage CP traffic) to the specified slot. • Enter the **slot** keyword and the slot number, as applicable.
Step 4	**service-policy input** *policy-map-name* **Example:** `Router(config-cp)# service-policy input control-plane-policy`	Attaches a QoS policy map to filter and manage CP traffic on a specified line card before the aggregate CP policy is applied. Note the following points: • **input** --Applies the specified policy map using the distributed switch engine to CP packets that are received from all interfaces on the line card. • *policy-map-name* --Name of a service policy map (created using the **policy-map** command) to be attached. The name can be a maximum of 40 alphanumeric characters. **Note** The **service-policy output** *policy-map-name* command is not supported for applying a QoS policy map for distributed control plane services.
Step 5	**end** **Example:** `Router(config-cp)# end`	(Optional) Returns to privileged EXEC mode.

Verifying Aggregate Control Plane Services

To display information about the service policy attached to the control plane for aggregate CP services, complete the following steps.

SUMMARY STEPS

1. **enable**
2. **show policy-map control-plane** [**all**] [**input** [**class** *class-name*] | output [class class-name]]
3. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.

	Command or Action	Purpose
Step 2	**show policy-map control-plane** [**all**] [**input** [**class** *class-name*] | output [class class-name]] **Example:** `Router# show policy-map control-plane all`	Displays information about the control plane. Note the following points: • **all** --(Optional) Service policy information about all QoS policies used in aggregate and distributed CP services. • **input** --(Optional) Statistics for the attached input policy. • **output** --(Optional) Statistics for the attached output policy. • **class** *class-name* --(Optional) Name of the traffic class whose configuration and statistics are displayed.
Step 3	**exit** **Example:** `Router(config-cp)# exit`	(Optional) Exits privileged EXEC mode.

Examples

The following example shows that the policy map TEST is associated with the control plane. This policy map polices traffic that matches the class map TEST, while allowing all other traffic (that matches the class map "class-default") to go through as is.

```
Router# show policy-map control-plane

Control Plane
Service-policy input:TEST
Class-map:TEST (match-all)
      20 packets, 11280 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match:access-group 101
      police:
        8000 bps, 1500 limit, 1500 extended limit
        conformed 15 packets, 6210 bytes; action:transmit
        exceeded 5 packets, 5070 bytes; action:drop
        violated 0 packets, 0 bytes; action:drop
        conformed 0 bps, exceed 0 bps, violate 0 bps
Class-map:class-default (match-any)
      105325 packets, 11415151 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match:any
```

Verifying Distributed Control Plane Services

To display information about the service policy attached to the control plane to perform distributed CP services, complete the following steps.

SUMMARY STEPS

1. **enable**
2. **show policy-map control-plane** [**all**][**slot** *slot-number*] [**input** [**class** *class-name*] | output [class class-name]]
3. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**show policy-map control-plane** [**all**][**slot** *slot-number*] [**input** [**class** *class-name*] | output [class class-name]] **Example:** `Router# show policy-map control-plane slot 2`	Displays information about the service policy used to apply distributed CP services on the router. Note the following points: • **all** --(Optional) Service policy information about all QoS policies used in aggregate and distributed CP services. • **slot** *slot-number* --(Optional) Service policy information about the QoS policy map used to perform distributed CP services on the specified line card. • **input** --(Optional) Statistics for the attached input policy map. • **output** --(Optional) Statistics for the attached output policy map. • **class** *class-name* --(Optional) Name of the traffic class whose configuration and statistics are displayed.
Step 3	**exit** **Example:** `Router# exit`	(Optional) Exits privileged EXEC mode.

Examples

The following example shows how to display information about the classes of CP traffic received from all interfaces on the line card in slot 1 to which the policy map TESTII is applied for distributed CP services. This policy map polices traffic that matches the traffic class TESTII, while allowing all other traffic (that matches the class map "class-default") to go through as is.

```
Router# show policy-map control-plane slot 1
Control Plane - slot 1
Service-policy input: TESTII (1048)
Class-map: TESTII (match-all) (1049/4)
        0 packets, 0 bytes
        5 minute offered rate 0 bps, drop rate 0 bps
        Match: protocol arp (1050)
        police:
            cir 8000 bps, bc 4470 bytes, be 4470 bytes
          conformed 0 packets, 0 bytes; actions:
            transmit
          exceeded 0 packets, 0 bytes; actions:
            drop
          violated 0 packets, 0 bytes; actions:
            drop
          conformed 0 bps, exceed 0 bps, violate 0 bps
Class-map: class-default (match-any) (1052/0)
        0 packets, 0 bytes
        5 minute offered rate 0 bps, drop rate 0 bps
        Match: any  (1053)
```

Configuration Examples for Control Plane Policing

This section contains examples that shows how to configure aggregate control plane services on both an input and an output interface:

Example Configuring Control Plane Policing on Input Telnet Traffic

The following example shows how to apply a QoS policy for aggregate CP services to Telnet traffic that is received on the control plane. Trusted hosts with source addresses 10.1.1.1 and 10.1.1.2 forward Telnet packets to the control plane without constraint, while allowing all remaining Telnet packets to be policed at the specified rate.

```
! Allow 10.1.1.1 trusted host traffic.
Router(config)# access-list 140 deny tcp host 10.1.1.1 any eq telnet

! Allow
10.1.1.2
 trusted host traffic.
Router(config)# access-list 140 deny tcp host 10.1.1.2 any eq telnet

! Rate-limit all other Telnet traffic.
Router(config)# access-list 140 permit tcp any any eq telnet
! Define class-map "telnet-class."
Router(config)# class-map telnet-class

Router(config-cmap)# match access-group 140
Router(config-cmap)# exit
Router(config)# policy-map control-plane-in
Router(config-pmap)# class telnet-class
Router(config-pmap-c)# police 80000 conform transmit exceed drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
! Define aggregate control plane service for the active route processor.
Router(config)# control-plane
Router(config-cp)# service-policy input control-plane-in
Router(config-cp)# end
```

Example Configuring Control Plane Policing on Output ICMP Traffic

The following example shows how to apply a QoS policy for aggregate CP services to Telnet traffic transmitted from the control plane. Trusted networks with source addresses 10.0.0.0 and 10.0.0.1 receive Internet Control Management Protocol (ICMP) port-unreachable responses without constraint, while allowing all remaining ICMP port-unreachable responses to be dropped:

```
! Allow
10.0.0.0
 trusted network traffic.
Router(config)# access-list 141 deny icmp 10.0.0.0 0.0.0.255 any port-unreachable

! Allow
10.0.0.1
 trusted network traffic.
Router(config)# access-list 141 deny icmp 10.0.0.1 0.0.0.255 any port-unreachable

! Rate-limit all other ICMP traffic.
Router(config)# access-list 141 permit icmp any any port-unreachable
Router(config)# class-map icmp-class
```

```
Router(config-cmap)# match access-group 141
Router(config-cmap)# exit
Router(config)# policy-map control-plane-out
! Drop all traffic that matches the class "icmp-class."
Router(config-pmap)# class icmp-class
Router(config-pmap-c)# drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# control-plane
! Define aggregate control plane service for the active route processor.
Router(config-cp)# service-policy output control-plane-out
Router(config-cp)# end
```

Additional References

The following sections provide references related to the Control Plane Policing feature.

Related Documents

Related Topic	Document Title
QoS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	*Cisco IOS Quality of Service Solutions Command Reference*
QoS features overview	"Quality of Service Overview" module
MQC	"Applying QoS Features Using the MQC" module
Security features overview	"Control Plane Security Overview" module in the *Cisco IOS Security Configuration Guide: Securing the Control Plane*
Control plane policing in Cisco IOS Release 12.2(18)SXD1 and later releases	For Catalyst 6500 series switches, see the "Configuring Control Plane Policing (CoPP)" module. For Cisco 7600 series routers, see the "Configuring Denial of Service Protection" module.
Enhanced RP protection	"ACL IP Options Selective Drop" module

Standards

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

MIBs

MIB	MIBs Link
• CISCO-CLASS-BASED-QOS-MIB **Note** Supported only in Cisco IOS Release 12.3(7)T.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator, found at the following URL: http://www.cisco.com/go/mibs

RFCs

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

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Control Plane Policing

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

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.

Table 5 Feature Information for Control Plane Policing

Feature Name	Releases	Feature Information
Control Plane Policing	12.2(18)S 12.3(4)T 12.3(7)T 12.0(29)S 12.2(18)SXD1 12.0(30)S 12.2(27)SBC 12.0(32)S 12.3(31)SB2 15.0(1)S	The Control Plane Policing feature allows users to configure a quality of service (QoS) filter that manages the traffic flow of control plane packets to protect the control plane of Cisco IOS routers and switches against reconnaissance and denial-of-service (DoS) attacks. For Release 12.2(18)S, this feature was introduced. For Release 12.3(4)T, this feature was integrated into Cisco IOS Release 12.3(4)T, and the output rate-limiting (silent mode operation) feature was added. For Release 12.3(7)T, the CISCO-CLASS-BASED-QOS-MIB was extended to manage control plane QoS policies, and the **police rate**command was introduced to support traffic policing on the basis of packets per second for control plane traffic. For Release 12.0(29)S, this feature was integrated into Cisco IOS Release 12.0(29)S. For Release 12.2(18)SXD1, this feature was integrated into Cisco IOS Release 12.2(18)SXD1. For Release 12.0(30)S, this feature was modified to include support for distributed control plane services on the Cisco 12000 series Internet router. For Release 12.2(27)SBC, this feature was integrated into Cisco IOS Release 12.2(27)SBC. For Release 12.0(32)S, this feature was modified to include support for aggregate control plane services on the Cisco 10720 Internet router. For Release 12.3(31)SB2, this feature was implemented on the

Feature Name	Releases	Feature Information
		Cisco 10000 series router for the PRE3. For Release 15.0(1)S, this feature was integrated into Cisco IOS Release 15.0(1)S.

Class-Based Policing

Feature History

Release	Modification
12.1(5)T	This command was introduced for Cisco IOS Release 12.1 T. A new Class-Based Policing algorithm was introduced. The **violate-action** option became available. This feature became available on Cisco 2600, 3600, 4500, 7200, and 7500 series routers.
12.2(2)T	The set-clp-transmit option for the *action* argument was added to the **police** command. The set-frde-transmit option for the *action* argument was added to the **police** command. The set-mpls-exp-transmit option for the *action* argument was added to the **police** command.
12.0(26)S	This feature was integrated into Cisco IOS Release 12.0(26)S for the Cisco 7200 and 7500 series routers. The name of the feature changed from *Traffic Policing* to *Class-Based Policing*.

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.

Feature Overview

This feature module describes the Class-Based Policing feature. The Class-Based Policing feature performs the following functions:

- Limits the input or output transmission rate of a class of traffic based on user-defined criteria
- Marks packets by setting the ATM Cell Loss Priority (CLP) bit, Frame Relay Discard Eligibility (DE) bit, IP precedence value, IP differentiated services code point (DSCP) value, MPLS experimental value, and Quality of Service (QoS) group.

Class-Based Policing allows you to control the maximum rate of traffic transmitted or received on an interface. The Class-Based Policing feature is applied when you attach a traffic policy contain the Class-Based Policing configuration to an interface. A traffic policy.

- Benefits, page 60
- Restrictions, page 61

Benefits

Bandwidth Management Through Rate Limiting

Class-Based Policing allows you to control the maximum rate of traffic transmitted or received on an interface. Class-Based Policing is often configured on interfaces at the edge of a network to limit traffic into or out of the network. In most Class-Based Policing configurations, traffic that falls within the rate parameters is transmitted, whereas traffic that exceeds the parameters is dropped or transmitted with a different priority.

Packet Marking

Packet marking allows you to partition your network into multiple priority levels or classes of service (CoS). A packet is marked and these markings can be used to identify and classify traffic for downstream devices. In some cases, such as ATM Cell Loss Priority (CLP) marking or Frame Relay Discard Eligibility (DE) marking, the marking is used to classify traffic.

- Use Class-Based Policing to set the IP precedence or DSCP values for packets entering the network. Networking devices within your network can then use the adjusted IP precedence values to determine how the traffic should be treated. For example, the Weighted Random Early Detection (WRED) feature uses the IP precedence values to determine the probability that a packet will be dropped.
- Use Class-Based Policing to assign packets to a QoS group. The router uses the QoS group to determine how to prioritize packets within the router.

Traffic can be marked without using the Class-Based Policing feature. If you want to mark traffic but do not want to use Class-Based Policing, see the "Marking Network Traffic" module.

Packet Prioritization for Frame Relay Frames

The Class-Based Policing feature allows users to mark the Frame Relay DE bit of the Frame Relay frame. The Frame Relay DE bit is one bit and, therefore, can be set to either 0 or 1. In congested environments, frames with the DE bit set to 1 are discarded before frames with the DE bit set to 0.

Packet Prioritization for ATM Cells

The Class-Based Policing feature allows users to mark the ATM CLP bit in ATM cells. The ATM CLP bit is used to prioritize packets in ATM networks. The ATM CLP bit is one bit and, therefore, can be set to either 0 or 1. In congested environments, cells with the ATM CLP bit set to 1 are discarded before cells with the ATM CLP bit set to 0.

Restrictions

- To use the *set-clp-transmit*action available with this feature, the Enhanced ATM Port Adapter (PA-A3) is required. Therefore, the *set-clp-transmit* action is not supported on any platform that does not support the PA-A3 adapter (such as the Cisco 2600 series router, the Cisco 3640 router, and the 4500 series router). For more information, see the documentation for your specific router.
- On a Cisco 7500 series router, Class-Based Policing can monitor Cisco Express Forwarding (CEF) switching paths only. In order to use the Class-Based Policing feature, Cisco Express Forwarding must be configured on both the interface receiving the packet and the interface sending the packet.
- On a Cisco 7500 series router, Class-Based Policing cannot be applied to packets that originated from or are destined to a router.
- Class-Based Policing can be configured on an interface or a subinterface.
- Class-Based Policing is not supported on the following interfaces:
  - Fast EtherChannel
  - Tunnel

Note Class-Based Policing is supported on tunnels that are using the Cisco generic routing encapsulation (GRE) tunneling protocol.

-
  - PRI
  - Any interface on a Cisco 7500 series router that does not support Cisco Express Forwarding

Prerequisites

On a Cisco 7500 series router, Cisco Express Forwarding (CEF) must be configured on the interface before Class-Based Policing can be used.

Configuration Tasks

- Configuring Traffic Policing, page 62
- Verifying Traffic Policing, page 62
- Troubleshooting Tips, page 62

Configuring Traffic Policing

Command	Purpose
Router(config-pmap-c)# **police** *bps burst-normal burst-max* **conform-action** *action* **exceed-action** *action* **violate-action** *action*	Specifies a maximum bandwidth usage by a traffic class. **Note** The Class-Based Policing feature works with a token bucket mechanism. There are currently two types of token bucket algorithms: a single token bucket algorithm and a two token bucket algorithm. A single token bucket system is used when the **violate-action** option is not specified, and a two token bucket system is used when the **violate-action** option is specified.

Verifying Traffic Policing

Use the **show policy-map interface** EXEC command to verify that the Class-Based Policing feature is configured on your interface. If the feature is configured on your interface, the **show policy-map interface** command output displays policing statistics:

```
Router# show policy-map interface
 Ethernet1/7
  service-policy output: x
    class-map: a (match-all)
      0 packets, 0 bytes
      5 minute rate 0 bps
      match: ip precedence 0
      police:
        1000000 bps, 10000 limit, 10000 extended limit
        conformed 0 packets, 0 bytes; action: transmit
        exceeded 0 packets, 0 bytes; action: drop
        conformed 0 bps, exceed 0 bps, violate 0 bps
```

Troubleshooting Tips

- Check the interface type. Verify that your interface is not mentioned in the nonsupported interface description in the Restrictions, page 61 section of this module.
- For input Class-Based Policing on a Cisco 7500 series router, verify that CEF is configured on the interface where Class-Based Policing is configured.
- For output Class-Based Policing on a Cisco 7500 series router, ensure that the incoming traffic is CEF-switched. Class-Based Policing cannot be used on the switching path unless CEF switching is enabled.

Monitoring and Maintaining Traffic Policing

Command	Purpose
Router# **show policy-map**	Displays all configured policy maps.

Command	Purpose
Router# **show policy-map** *policy-map-name*	Displays the user-specified policy map.
Router# **show policy-map interface**	Displays statistics and configurations of all input and output policies that are attached to an interface.
Router# **show policy-map interface service instance**	Displays the policy map information for a given service instance under a port channel.

Configuration Examples

Example Configuring a Service Policy that Includes Traffic Policing

In the following example, Class-Based Policing is configured with the average rate at 8000 bits per second, the normal burst size at 1000 bytes, and the excess burst size at 1000 bytes for all packets leaving Fast Ethernet interface 0/0.

```
class-map access-match
match access-group 1
exit
policy-map police-setting
class access-match
police 8000 1000 1000 conform-action transmit exceed-action set-qos-transmit 1 violate-
action drop
exit
exit
service-policy output police-setting
```

The treatment of a series of packets leaving Fast Ethernet interface 0/0 depends on the size of the packet and the number of bytes remaining in the conform and exceed token buckets. The series of packets are policed based on the following rules:

- If the previous arrival of the packet was at T1 and the current arrival of the packet is at T, the bucket is updated with T -T1 worth of bits based on the token arrival rate. The refill tokens are placed in the conform bucket. If the tokens overflow the conform bucket, the overflow tokens are placed in the exceed bucket. The token arrival rate is calculated as follows:

(time between packets <which is equal to T - T1> * policer rate)/8 bytes

- If the number of bytes in the conform bucket B is greater than or equal to 0, the packet conforms and the conform action is taken on the packet. If the packet conforms, B bytes are removed from the conform bucket and the conform action is taken. The exceed bucket is unaffected in this scenario.
- If the number of bytes in the conform bucket B is less than 0, the excess token bucket is checked for bytes by the packet. If the number of bytes in the exceed bucket B is greater than or equal to 0, the exceed action is taken and B bytes are removed from the exceed token bucket. No bytes are removed from the conform bucket in this scenario.
- If the number bytes in the exceed bucket B is fewer than 0, the packet violates the rate and the violate action is taken. The action is complete for the packet.

In this example, the initial token buckets starts full at 1000 bytes. If a 450-byte packet arrives, the packet conforms because enough bytes are available in the conform token bucket. The conform action (send) is taken by the packet and 450 bytes are removed from the conform token bucket (leaving 550 bytes).

If the next packet arrives 0.25 seconds later, 250 bytes are added to the conform token bucket

((0.25 * 8000)/8), leaving 800 bytes in the conform token bucket. If the next packet is 900 bytes, the packet does not conform because only 800 bytes are available in the conform token bucket.

The exceed token bucket, which starts full at 1000 bytes (as specified by the excess burst size) is then checked for available bytes. Because enough bytes are available in the exceed token bucket, the exceed action (set the QoS transmit value of 1) is taken and 900 bytes are taken from the exceed bucket (leaving 100 bytes in the exceed token bucket.

If the next packet arrives 0.40 seconds later, 400 bytes are added to the token buckets ((.40 * 8000)/8). Therefore, the conform token bucket now has 1000 bytes (the maximum number of tokens available in the conform bucket) and 200 bytes overflow the conform token bucket (because it only 200 bytes were needed to fill the conform token bucket to capacity). These overflow bytes are placed in the exceed token bucket, giving the exceed token bucket 300 bytes.

If the arriving packet is 1000 bytes, the packet conforms because enough bytes are available in the conform token bucket. The conform action (transmit) is taken by the packet, and 1000 bytes are removed from the conform token bucket (leaving 0 bytes).

If the next packet arrives 0.20 seconds later, 200 bytes are added to the token bucket ((.20 * 8000)/8). Therefore, the conform bucket now has 200 bytes. If the arriving packet is 400 bytes, the packet does not conform because only 200 bytes are available in the conform bucket. Similarly, the packet does not exceed because only 300 bytes are available in the exceed bucket. Therefore, the packet violates and the violate action (drop) is taken.

Additional References

The following sections provide references related to Traffic Policing.

Related Documents

Related Topic	Document Title
Traffic policing	"Traffic Policing" module
Modular Quality of Service Command-Line Interface (MQC)	"Applying QoS Features Using the MQC" module

Standards

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

MIBs

MIBs	MIBs Link
Class-Based Quality of Service MIB • CISCO-CLASS-BASED-QOS-MIB • CISCO-CLASS-BASED-QOS-CAPABILITY-MIB	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFCs	Title
RFC 2697	*A Single Rate Three Color Marker*

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

QoS Percentage-Based Policing

The QoS Percentage-Based Policing feature allows you to configure traffic policing and traffic shaping on the basis of a percentage of bandwidth available on the interface. This feature also allows you to specify the committed burst (bc) size and the excess burst (be) size (used for configuring traffic policing) in milliseconds (ms). Configuring traffic policing in this manner enables you to use the same policy map for multiple interfaces with differing amounts of bandwidth.

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.

Prerequisites for QoS Percentage-Based Policing

- For input traffic policing on a Cisco 7500 series router, verify that distributed Cisco Express Forwarding (dCEF) is enabled on the interface on which traffic policing is configured.
- For output traffic policing on a Cisco 7500 series router, ensure that the incoming traffic is dCEF-switched. Traffic policing cannot be used on the switching path unless dCEF switching is enabled.

Restrictions for QoS Percentage-Based Policing

The **shape** (percent) command, when used in "child" (nested) policy maps, is not supported on the Cisco 7500, the Cisco 7200, or lower series routers. Therefore, the **shape** (percent) command cannot be configured for use in nested policy maps on these routers.

Information About QoS Percentage-Based Policing

Benefits for QoS Percentage-Based Policing

This feature provides the ability to configure traffic policing and traffic shaping on the basis of a percentage of bandwidth available on an interface, and it allows you to specify burst sizes in milliseconds. Configuring traffic policing and traffic shaping in this manner enables you to use the same policy map for multiple interfaces with differing amounts of bandwidth. That is, you do not have to recalculate the bandwidth for each interface or configure a different policy map for each type of interface.

Defining Class and Policy Maps for QoS Percentage-Based Policing

To configure the QoS Percentage-Based Policing feature, you must define a traffic class, configure a policy map, and then attach that policy map to the appropriate interface. These three tasks can be accomplished by using the Modular Quality of Service (QoS) Command-Line Interface (CLI) (MQC).

The MQC is a command-line interface that allows you to define traffic classes, create and configure traffic policies (policy maps), and then attach these traffic policies to interfaces.

In the MQC, the **class-map** command is used to define a traffic class (which is then associated with a traffic policy). The purpose of a traffic class is to classify traffic.

The MQC consists of the following three processes:

- Defining a traffic class with the **class-map** command.
- Creating a traffic policy by associating the traffic class with one or more QoS features (using the **policy-map** command).
- Attaching the traffic policy to the interface with the **service-policy** command.

A traffic class contains three major elements: a name, a series of match commands, and, if more than one **match** command exists in the traffic class, an instruction on how to evaluate these **match** commands (that is, match-all or match-any). The traffic class is named in the **class-map** command line; for example, if you enter the **class-map cisco** command while configuring the traffic class in the CLI, the traffic class would be named "cisco".

The **match** commands are used to specify various criteria for classifying packets. Packets are checked to determine whether they match the criteria specified in the **match** commands. If a packet matches the specified criteria, that packet is considered a member of the class and is forwarded according to the QoS specifications set in the traffic policy. Packets that fail to meet any of the matching criteria are classified as members of the default traffic class.

Traffic Regulation Mechanisms and Bandwidth Percentages

Cisco IOS quality of service (QoS) offers two kinds of traffic regulation mechanisms--traffic policing and traffic shaping. A traffic policer typically drops traffic that violates a specific rate. A traffic shaper typically delays excess traffic using a buffer to hold packets and shapes the flow when the data rate to a queue is higher than expected.

Traffic shaping and traffic policing can work in tandem and can be configured in a class map. Class maps organize data packets into specific categories ("classes") that can, in turn, receive a user-defined QoS treatment when used in policy maps (sometimes referred to as "service policies").

Before this feature, traffic policing and traffic shaping were configured on the basis of a user-specified amount of bandwidth available on the interface. Policy maps were then configured on the basis of that specific amount of bandwidth, meaning that separate policy maps were required for each interface.

This feature provides the ability to configure traffic policing and traffic shaping on the basis of a *percentage* of bandwidth available on the interface. Configuring traffic policing and traffic shaping in this manner enables customers to use the same policy map for multiple interfaces with differing amounts of bandwidth.

Configuring traffic policing and shaping on the basis of a percentage of bandwidth is accomplished by using the **police** (percent) and **shape** (percent) commands.

Burst Size in Milliseconds Option

The purpose of the burst parameters (bc and be) is to drop packets gradually, as is done with Weighted Random Early Detection (WRED), and to avoid tail drop. Setting sufficiently high burst values helps to ensure good throughput.

This feature allows you the option of specifying the committed burst (bc) size and the extended burst (be) as milliseconds (ms) of the class bandwidth when you configure traffic policing. The number of milliseconds is used to calculate the number of bytes that will be used by the QoS Percentage-Based Policing feature.

Specifying these burst sizes in milliseconds is accomplished by using the **bc** and **be** keywords (and their associated arguments) of the **police** (percent) and **shape** (percent) commands.

How to Configure QoS Percentage-Based Policing

Configuring a Class and Policy Map for Percentage-Based Policing

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **policy-map** *policy-name*
4. **class** {*class-name* **class-default**}
5. **police cir percent** *percentage* [*burst-in-ms*] [**bc** *conform-burst-in-msec* **ms**] [**be** *peak-burst-in-msec* **ms**] [**pir percent** *percent*]
6. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**policy-map** *policy-name* **Example:** `Router(config)# policy-map policy1`	Specifies the name of the policy map to be created. Enters policy-map configuration mode. • Enter the policy map name.
Step 4	**class** {*class-name* **class-default**} **Example:** `Router(config-pmap)# class class1`	Specifies the class so that you can configure or modify its policy. Enters policy-map class configuration mode. • Enter the class name or specify the default class (class-default).

	Command or Action	Purpose
Step 5	**police cir percent** *percentage* [*burst-in-ms*] [**bc** *conform-burst-in-msec* **ms**] [**be** *peak-burst-in-msec* **ms**] [**pir percent** *percent*] **Example:** `Router(config-pmap-c)# police cir percent 20 bc 300 ms be 400 ms pir percent 40`	Configures traffic policing on the basis of the specified bandwidth percentage and optional burst sizes. Enters policy-map class police configuration mode. • Enter the bandwidth percentage and optional burst sizes.
Step 6	**exit** **Example:** `Router(config-pmap-c-police)# exit`	Exits policy-map class police configuration mode.

Attaching the Policy Map to an Interface for Percentage-Based Policing

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **pvc** [*name*] *vpi* / *vci* [**ilmi** | **qsaal** | **smds**]
5. **service-policy** {**input**| **output**} *policy-map-name*
6. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose

Step 3	**interface** *type number*
	
	Example:
	`Router(config)#`
	`interface serial4/0`
		Configures an interface (or subinterface) type and enters interface configuration mode.
		• Enter the interface type number.
		Note Depending on the needs of your network, you may need to attach the policy map to a subinterface, an ATM PVC, a Frame Relay DLCI, or other type of interface.

Step 4	**pvc** [*name*] *vpi* / *vci* [**ilmi** | **qsaal** | **smds**] **Example:** `Router(config-if)# pvc cisco 0/16 ilmi`	(Optional) Creates or assigns a name to an ATM PVC and specifies the encapsulation type on an ATM PVC. Enters ATM VC configuration mode. **Note** This step is required only if you are attaching the policy map to an ATM PVC. If you are not attaching the policy map to an ATM PVC, skip this step and proceed with Attaching the Policy Map to an Interface for Percentage-Based Policing, page 71.

Step 5	**service-policy** {**input**| **output**} *policy-map-name*
	
	Example:
	`Router(config-if)#`
	`service-policy input policy1`
	
	Example:
		Specifies the name of the policy map to be attached to the input *or* output direction of the interface.
		Note Policy maps can be configured on ingress or egress routers. They can also be attached in the input or output direction of an interface. The direction (input or output) and the router (ingress or egress) to which the policy map should be attached varies according your network configuration. When using the **service-policy** command to attach the policy map to an interface, be sure to choose the router and the interface direction that are appropriate for your network configuration.
		• Enter the policy map name.

Step 6	**exit** **Example:** `Router(config-if)# exit`	(Optional) Exits interface configuration mode.

Verifying the Percentage-Based Policing Configuration

SUMMARY STEPS

1. **enable**
2. **show class-map** [*class-map-name*]
3. **show policy-map interface** *interface-name*
4. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**show class-map** [*class-map-name*] **Example:** `Router# show class-map class1`	Displays all information about a class map, including the match criterion. • Enter class map name.

Step 3	**show policy-map interface** *interface-name*
	
	Example:
	`Router#`
	`show policy-map interface serial4/0`
		Displays the packet statistics of all classes that are configured for all service policies either on the specified interface or subinterface or on a specific PVC on the interface.
		• Enter the interface name.

Step 4	**exit** **Example:** `Router# exit`	(Optional) Exits privileged EXEC mode.

• Troubleshooting Tips for Percentage-Based Policing, page 73

Troubleshooting Tips for Percentage-Based Policing

The commands in the Verifying the Percentage-Based Policing Configuration, page 72 section allow you to verify that you achieved the intended configuration and that the feature is functioning correctly. If, after using the **show** commands listed above, you find that the configuration is not correct or the feature is not functioning as expected, perform these operations:

If the configuration is not the one you intended, complete the following procedures:

1. Use the **show running-config** command and analyze the output of the command.
2. If the policy map does not appear in the output of the **show running-config** command, enable the **logging console** command.
3. Attach the policy map to the interface again.

If the packets are not being matched correctly (for example, the packet counters are not incrementing correctly), complete the following procedures:

1. Run the **show policy-map**command and analyze the output of the command.
2. Run the **show running-config** command and analyze the output of the command.

3 Use the **show policy-map interface** command and analyze the output of the command. Check the the following findings:

   a If a policy map applies queueing, and the packets are matching the correct class, but you see unexpected results, compare the number of the packets in the queue with the number of the packets matched.

   b If the interface is congested, and only a small number of the packets are being matched, check the tuning of the transmission (tx) ring, and evaluate whether the queueing is happening on the tx ring. To do this, use the **show controllers** command, and look at the value of the tx count in the output of the command.

Configuration Examples for QoS Percentage-Based Policing

Specifying Traffic Policing on the Basis of a Bandwidth Percentage Example

The following example configures traffic policing using a committed information rate (CIR) and a peak information rate (PIR) on the basis of a percentage of bandwidth. In this example, a CIR of 20 percent and a PIR of 40 percent have been specified. Additionally, an optional bc value and be value (300 ms and 400 ms, respectively) have been specified.

```
Router> enable
Router# configure terminal
Router(config)# policy-map policy1
Router(config-pmap)# class class1
Router(config-pmap-c)# police cir percent 20 bc 300 ms be 400 ms pir percent 40

Router(config-pmap-c-police)# exit
```

After the policy map and class maps are configured, the policy map is attached to interface as shown in the following example.

```
Router> enable
Router# configure terminal
Router(config-if)#

interface serial4/0
Router(config-if)#

service-policy input policy1
Router(config-if)# exit
```

Verifying the Percentage-Based Policing Configuration Example

This section contains sample output from the **show policy-map interface** command and the **show policy-map** command. The output from these commands can be used to verify and monitor the feature configuration on your network.

The following is sample output from the **show policy-map** command. This sample output displays the contents of a policy map called "policy1." In policy 1, traffic policing on the basis of a CIR of 20 percent

has been configured, and the bc and be have been specified in milliseconds. As part of the traffic policing configuration, optional conform, exceed, and violate actions have been specified.

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

The following is sample output from the **show policy-map interface** command. This sample displays the statistics for the serial 2/0 interface on which traffic policing has been enabled. The committed burst (bc) and excess burst (be) are specified in milliseconds (ms).

```
Router# show policy-map interface serial2/0
 Serial2/0
  Service-policy output: policy1 (1050)
    Class-map: class1 (match-all) (1051/1)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: ip precedence 0  (1052)
      police:
          cir 20 % bc 300 ms
          cir 409500 bps, bc 15360 bytes
          pir 40 % be 400 ms
          pir 819000 bps, be 40960 bytes
        conformed 0 packets, 0 bytes; actions:
          transmit
        exceeded 0 packets, 0 bytes; actions:
          drop
        violated 0 packets, 0 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps, violate 0 bps
    Class-map: class-default (match-any) (1054/0)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any  (1055)
        0 packets, 0 bytes
        5 minute rate 0 bps
```

In this example, the CIR and PIR are displayed in bps, and both the committed burst (bc) and excess burst (be) are displayed in bits.

The CIR, PIR bc, and be are calculated on the basis of the formulas described below.

Formula for Calculating the CIR

When calculating the CIR, the following formula is used:

- CIR percentage specified (as shown in the output of the **show policy-map**command) * bandwidth (BW) of the interface (as shown in the output of the **show interfaces** command) = total bits per second

On serial interface 2/0, the bandwidth (BW) is 2048 kbps. To see the bandwidth of the interface, use the **show interfaces** command. A sample is shown below:

```
Router# show interfaces serial2/0
Serial2/0 is administratively down, line protocol is down
  Hardware is M4T
  MTU 1500 bytes, BW 2048 Kbit, DLY 20000 usec, rely 255/255, load 1/255
```

The following values are used for calculating the CI:.

20 % * 2048 kbps = 409600 bps

Formula for Calculating the PIR

When calculating the PIR, the following formula is used:

- PIR percentage specified (as shown in the output of the **show policy-map**command) * bandwidth (BW) of the interface (as shown in the output of the **show interfaces** command) = total bits per second

On serial interface 2/0, the bandwidth (BW) is 2048 kbps. To see the bandwidth of the interface, use the **show interfaces** command. A sample is shown below:

```
Router# show interfaces serial2/0
Serial2/0 is administratively down, line protocol is down
  Hardware is M4T
  MTU 1500 bytes, BW 2048 Kbit, DLY 20000 usec, rely 255/255, load 1/255
```

The following values are used for calculating the PIR:

40 % * 2048 kbps = 819200 bps

Note Discrepancies between this total and the total shown in the output of the **show policy-map interface** command can be attributed to a rounding calculation or to differences associated with the specific interface configuration.

Formula for Calculating the Committed Burst (bc)

When calculating the bc, the following formula is used:

- The bc in milliseconds (as shown in the **show policy-map** command) * the CIR in bits per seconds = total number bytes

The following values are used for calculating the bc:

300 ms * 409600 bps = 15360 bytes

Formula for Calculating the Excess Burst (be)

When calculating the bc and the be, the following formula is used:

- The be in milliseconds (as shown in the **show policy-map** command) * the PIR in bits per seconds = total number bytes

The following values are used for calculating the be:

400 ms * 819200 bps = 40960 bytes

Additional References

The following sections provide references related to the QoS Percentage-Based Policing feature.

Related Documents

Related Topic	Document Title
QoS commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	*Cisco IOS Quality of Service Solutions Command Reference*
Modular QoS Command-Line Interface (CLI) (MQC), including information about attaching policy maps	"Applying QoS Features Using the MQC" module
Traffic shaping and traffic policing	"Policing and Shaping Overview" module
Commands related to dCEF	*Cisco IOS Switching Command Reference*

Standard

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

MIB

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

RFCs

RFC	Title
RFC 2697	*A Single Rate Three Color Marker*
RFC 2698	*A Two Rate Three Color Marker*

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for QoS Percentage-Based Policing

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

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.

Table 6 Feature Information for QoS Percentage-Based Policing

Feature Name	Releases	Feature Information
QoS Percentage-Based Policing	12.2(13)T 12.0(28)S 12.2(28)SB 15.0(1)S	The QoS Percentage-Based Policing feature allows you to configure traffic policing and traffic shaping on the basis of a *percentage* of bandwidth available on the interface. This feature also allows you to specify the committed burst (bc) size and the excess burst (be) size (used for configuring traffic policing) in milliseconds (ms). Configuring traffic policing in this manner enables you to use the same policy map for multiple interfaces with differing amounts of bandwidth. In Release 12.2(13)T, this feature was introduced. In Release 12.0(28)S, the option of specifying committed (conform) burst (bc) and excess (peak) burst (be) sizes in milliseconds was added. In Release 12.2(28)SB, this feature was integrated in Cisco IOS Release 12.2(28)SB. In Release 15.0(1)S, this feature was integrated in Cisco IOS Release 15.0(1)S. The following commands were introduced or modified: **police (percent)**, **shape (percent)**, **show policy-map**, **show policy-map interface**.

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.