



JunosE™ Software for E Series™ Broadband Services Routers

Rate-Limit Profile Management

Release
12.3.x



Published: 2011-10-14

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

JunosE™ Software for E Series™ Broadband Services Routers Rate-Limit Profile Management
Release 12.3.x
Copyright © 2011, Juniper Networks, Inc.
All rights reserved.

Revision History
October 2011—FRS JunosE 12.3.x

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	ix
	E Series and JunosE Documentation and Release Notes	ix
	Audience	ix
	E Series and JunosE Text and Syntax Conventions	ix
	Obtaining Documentation	xi
	Documentation Feedback	xi
	Requesting Technical Support	xi
	Self-Help Online Tools and Resources	xii
	Opening a Case with JTAC	xii
Part 1	Overview	
Chapter 1	How Rate-Limit Profiles Work	3
	Rate Limits for Interfaces Overview	3
	Hierarchical Rate Limits Overview	4
	Hierarchical Classifier Groups	5
	Hierarchical Rate-Limit Profiles	6
	Hierarchical Rate-Limit Actions	6
	Example: Multiple Flows Sharing Preferred Bandwidth Rate-Limiting Hierarchical Policy	8
	Example: Multiple Flows Sharing a Rate Limit Hierarchical Policy	9
	Example: Shared Pool of Additional Bandwidth with Select Flows Rate-Limiting Hierarchical Policy	10
	Example: Aggregate Marking with Oversubscription Rate-Limiting Hierarchical Policy	12
	Color-Aware Configuration for Rate-Limiting Hierarchical Policy	14
	Percent-Based Rates for Rate-Limit Profiles Overview	15
	Policy Parameter Reference-Rate	15
	Specifying Rates Within Rate-Limit Profiles	16
	Specifying Burst Sizes	17
	Using Service Manager with Merged Policies	17
	Policy Parameter Configuration Considerations	17
	One-Rate Rate-Limit Profiles Overview	19
	Two-Rate Rate-Limits Overview	20
	Bandwidth Management Overview	22
	Examples: One-Rate Rate-Limit Profile	23
	Examples: Two-Rate Rate-Limit Profile	24
	Examples: Rate-Limiting Individual or Aggregate Packet Flows	24

Part 2	Configuration	
Chapter 2	Configuration Tasks for Rate-Limit Profile Management	29
	Creating Rate-Limit Profiles	29
	Policy Parameter Quick Configuration	34
	Setting the Committed Action for a Rate-Limit Profile	34
	Setting the Committed Burst for a Rate-Limit Profile	35
	Setting the Committed Rate for a Rate-Limit Profile	36
	Setting the Conformed Action for a Rate-Limit Profile	36
	Setting the Exceeded Action for a Rate-Limit Profile	36
	Setting the Excess Burst for a Rate-Limit Profile	37
	Setting the Mask Value for MPLS Rate-Limit Profiles	37
	Setting the Mask Value for IP and IPv6 Rate-Limit Profiles	37
	Setting the Peak Rate for Rate-Limit Profiles	37
	Rate-Limiting Traffic Flows	38
Chapter 3	Configuration Tasks for One-Rate Rate-Limit Profile Management	39
	Creating a One-Rate Rate-Limit Profile	39
	Configuring a TCP-Friendly One-Rate Rate-Limit Profile	40
	Setting a One-Rate Rate-Limit Profile	42
Chapter 4	Configuration Tasks for Two-Rate Rate-Limit Profile Management	45
	Creating a Two-Rate Rate-Limit Profile	45
	Setting the Peak Burst for Two-Rate Rate-Limit Profiles	46
	Setting a Two-Rate Rate-Limit-Profile	46
Part 3	Administration	
Chapter 5	Monitoring Tasks	51
	Monitoring Rate-Limit Profiles	51
	Monitoring Control Plane Policer Information	53
Part 4	Index	
	Index	57

List of Figures

Part 1	Overview	
Chapter 1	How Rate-Limit Profiles Work	3
	Figure 1: Multiple Flows Sharing Preferred Bandwidth	8
	Figure 2: Multiple Packet Flows Sharing a Rate Limit	9
	Figure 3: Shared Pool of Additional Bandwidth with Select Flows	11
	Figure 4: Aggregate Marking with Oversubscription	13
	Figure 5: Congestion Management	23

List of Tables

	About the Documentation	ix
	Table 1: Notice Icons	x
	Table 2: Text and Syntax Conventions	x
Part 1	Overview	
Chapter 1	How Rate-Limit Profiles Work	3
	Table 3: Policy Action Applied Based on Rate Settings and Traffic Rate	21
	Table 4: Two-Rate Rate-Limit Profile Algorithms	22
Part 2	Configuration	
Chapter 3	Configuration Tasks for One-Rate Rate-Limit Profile Management	39
	Table 5: TCP-Friendly One-Rate Rate-Limit Profile Algorithms	41
	Table 6: One-Rate Rate-Limit-Profile Defaults	42
Chapter 4	Configuration Tasks for Two-Rate Rate-Limit Profile Management	45
	Table 7: Two-Rate Rate-Limit-Profile Defaults	47
Part 3	Administration	
Chapter 5	Monitoring Tasks	51
	Table 8: show rate-limit-profile Output Fields	52
	Table 9: show control-plane policer Output Fields	53

About the Documentation

- E Series and JunosE Documentation and Release Notes on page ix
- Audience on page ix
- E Series and JunosE Text and Syntax Conventions on page ix
- Obtaining Documentation on page xi
- Documentation Feedback on page xi
- Requesting Technical Support on page xi

E Series and JunosE Documentation and Release Notes

For a list of related JunosE documentation, see
<http://www.juniper.net/techpubs/software/index.html> .

If the information in the latest release notes differs from the information in the documentation, follow the *JunosE Release Notes*.

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at
<http://www.juniper.net/techpubs/> .

Audience

This guide is intended for experienced system and network specialists working with Juniper Networks E Series Broadband Services Routers in an Internet access environment.

E Series and JunosE Text and Syntax Conventions

Table 1 on page x defines notice icons used in this documentation.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2 on page x defines text and syntax conventions that we use throughout the E Series and JunosE documentation.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents commands and keywords in text.	<ul style="list-style-type: none"> Issue the clock source command. Specify the keyword exp-msg.
Bold text like this	Represents text that the user must type.	host1(config)#traffic class low-loss1
Fixed-width text like this	Represents information as displayed on your terminal's screen.	host1#show ip ospf 2 Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an Area Border Router (ABR)
<i>Italic text like this</i>	<ul style="list-style-type: none"> Emphasizes words. Identifies variables. Identifies chapter, appendix, and book names. 	<ul style="list-style-type: none"> There are two levels of access: <i>user</i> and <i>privileged</i>. <i>clusterId</i>, <i>ipAddress</i>. <i>Appendix A, System Specifications</i>
Plus sign (+) linking key names	Indicates that you must press two or more keys simultaneously.	Press Ctrl + b.
Syntax Conventions in the Command Reference Guide		
Plain text like this	Represents keywords.	terminal length
<i>Italic text like this</i>	Represents variables.	<i>mask</i> , <i>accessListName</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
(pipe symbol)	Represents a choice to select one keyword or variable to the left or to the right of this symbol. (The keyword or variable can be either optional or required.)	diagnostic line
[] (brackets)	Represent optional keywords or variables.	[internal external]
[]* (brackets and asterisk)	Represent optional keywords or variables that can be entered more than once.	[level1 level2 l1]*
{ } (braces)	Represent required keywords or variables.	{ permit deny } { in out } { clusterId ipAddress }

Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the Technical Documentation page on the Juniper Networks Web site at <http://www.juniper.net/>.

To download complete sets of technical documentation to create your own documentation CD-ROMs or DVD-ROMs, see the Portable Libraries page at

<http://www.juniper.net/techpubs/resources/index.html>

Copies of the Management Information Bases (MIBs) for a particular software release are available for download in the software image bundle from the Juniper Networks Web site at <http://www.juniper.net/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation to better meet your needs. Send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract,

or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf> .
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/> .
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/> .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html> .

PART 1

Overview

- [How Rate-Limit Profiles Work on page 3](#)

CHAPTER 1

How Rate-Limit Profiles Work

- [Rate Limits for Interfaces Overview on page 3](#)
- [Hierarchical Rate Limits Overview on page 4](#)
- [Percent-Based Rates for Rate-Limit Profiles Overview on page 15](#)
- [One-Rate Rate-Limit Profiles Overview on page 19](#)
- [Two-Rate Rate-Limits Overview on page 20](#)
- [Bandwidth Management Overview on page 22](#)

Rate Limits for Interfaces Overview

To configure rate limiting for interfaces, you first create a rate-limit profile, which is a set of bandwidth attributes and associated actions. Your router supports two types of rate-limit profiles—one-rate and two-rate—for IP, IPv6, LT2P, and MPLS Layer 2 transport traffic. You next create a policy list with a rule that has rate limit as the action and associate a rate-limit profile with this rule.

You configure rate limit profiles from Global Configuration Mode.



NOTE: Commands that you issue in Rate Limit Profile Configuration mode do not take effect until you exit from that mode.

When packets enter an interface that has a rate-limit profile applied, the router performs the following:

- Counts the number of bytes (packets) over time
- Categorizes each packet as committed, conformed, or exceeded
- Assigns a transmit, drop, or mark action



NOTE: Mark actions and mark values are supported only on IP, IPv6, and MPLS rate-limit profiles. They are not supported on hierarchical rate limits, but are replaced by color-mark profiles.

An additional function of rate limiting is to apply a color code to packets assigned to each category: green for committed, yellow for conformed, and red for exceeded. The system uses the color code internally to indicate drop preference when an outbound interface is congested.

Rate limiters are implemented using a dual token bucket scheme: a token bucket for conformed (yellow) packets and a token bucket for committed (green) packets. One token is synonymous with one byte. The capacity of the buckets is the maximum number of tokens that can be placed in each bucket.

You configure the bucket capacity with the peak burst parameter or the committed burst parameter. The burst parameters are in bytes (not bytes per second), which is the number of tokens in a full bucket. When a packet passes through a rate limiter, its size is compared to the contents of both buckets, the packet is categorized, and the rate-limiter action is taken on the packet.

Peak rate and committed rate determine the fill rate of their respective buckets. If you set the committed rate to 128,000 bps, tokens are added to the committed (green) bucket at a rate of 128,000 bps (16 K bytes per second), regardless of the traffic. If no traffic passes through the rate limiter, the bucket continues to fill until it reaches the committed burst setting.

Traffic passes through the rate limiter causing a draining of tokens. The drain rate is dependent on how large the packets are and how much time elapses between packets. At any given instant the level of tokens in each bucket is a function of the fill rate, size of packets, and elapsed time between packets.

When packets are received on an interface with a rate limiter applied, the level of tokens in each bucket dynamically changes in both of the following ways:

- Tokens are added every 100-ms sample period
- Tokens are removed based on the size and rate of incoming packets

Hierarchical Rate Limits Overview

In another type of rate limiting, rate-limit hierarchies enable lower priority traffic to access unused bandwidth allocated for real-time traffic, such as voice or video, during times when no real-time traffic is flowing. IP subscribers receive multiple services, such as Web, video, and file transfer, that have a maximum bandwidth. A rate-limit hierarchy can apply a common rate limit to several classified flows, enabling them to share bandwidth according to the preferences set in the hierarchical rate limits.

You can also use rate-limit hierarchies in a layer 2 (ATM) access network for DSL where many routing gateways lead into one Broadband Access Server. The Broadband Access Server uses rate-limit hierarchies to allocate shareable bandwidth to each routing gateway, which enables unused bandwidth from one routing gateway to be used by others. The hierarchy in the rate limit represents the hierarchy in the access network.

Rate-limit hierarchies enable you to share unused bandwidth dynamically, taking unused preferred bandwidth. They also enable real-time traffic to use all guaranteed bandwidth

at any time without violating the configured limit on the total interface bandwidth. While preferred traffic fluctuates, the interface rate limit adjusts, dropping non-preferred packets to keep the total flow through the interface under a configured maximum rate, because preferred packets cannot be dropped by the shared rate limits, only by their individual rate limits.

Shared rate limits in the hierarchy keep the combined traffic below a configured maximum without dropping preferred packets. Preferred packets always reduce tokens on these rate limits, making their token counts negative, if necessary. Later non-preferred packets are then dropped in greater volume, bringing the total traffic through the shared rate limit below its configured maximum.

Every packet passing through a rate limit hierarchy has an owner, which is the last rate limit that can modify the packet; for example, by changing its color or dropping it. Preferred packets are owned by their individual preferred rate limits, which do not transfer ownership of the packet while the packet traverses the hierarchy. Ownership of non-preferred packets is transferred while they move from one rate-limit to the next in the hierarchy, so shared rate limits can change the packet color or drop them.

Hierarchical Classifier Groups

Rate-limit hierarchies can be intra-interface, where different flows from classifier groups are in one policy attachment on an interface. Each time the policy is attached to another interface the rate-limit hierarchy is replicated, with no rate limits shared between attachments. Hierarchical rate-limits are only applied at forwarding interfaces because they provide the most accurate classification of packets.

You can configure rate-limit hierarchies by defining a hierarchy of policy classifier and parent groups, each with a rate limit. This hierarchy applies to the packet flow on one interface attachment for the policy. Each policy attachment creates its own copy of the rate-limit hierarchy. There are no shared rate limits across interface attachments.

A policy-based rate-limit hierarchy consists of classifier groups with an aggregate node policy object. Aggregate nodes create the interior nodes of a policy-based hierarchy; they are not classifier groups and the only policy rule applicable to them is the rate limit rule. Every classifier group or aggregate node can select another aggregate node as its parent. The policy manager ensures that these choices always result in a hierarchy. Not every classifier group with a parent aggregate node must have a rate limit rule; multiple classifier groups can share a common parent group, which may have a rate limit rule.

A policy imposes a limit of three parent groups that can be traversed from any classifier group. However, the total number of parent groups in one policy can be up to 512, but every packet must pass through no more than three parent groups at any point.

In a hierarchy of rate limits, a rate limit can be color-blind or color-aware; color-blind rate limits run the same algorithm for all packets, regardless of their color. Color-aware rate limits can change the algorithm used, depending on the color of the incoming packet (possibly set in the previous rate limit or an earlier policy, such as a VLAN policy on ingress or an IP policy). The color mark profile action changes the ToS field for the packet, depending on packet type (EXP for MPLS, DSCP or ToS for IPv4), and transmits the

packet. If the mark action uses a color-mark profile, the ToS values marked can depend on the color of the packet.

Hierarchical Rate-Limit Profiles

Hierarchical rate-limit profiles are independent from interface types. You can apply the green, yellow, or red mark values to the rate-limit profile for every type of forwarding interface that accepts ToS marking for packets. The same rate limit can be reused for a different interface type. Hierarchical rate limits have two-rate or TCP-friendly rate types.

The value applied to the ToS field is configured in the CLACL group for green, yellow, or red packets but the coloring of the packet as green, yellow, or red depends on the entire rate-limit hierarchy.

- Preferred packets are transmitted unconditionally. Rate limits that process packets transmitted unconditionally always decrement their token count, if necessary, making it negative.
- Red packets cannot be transmitted unconditionally, to avoid cases where an aggregate rate limit is oversubscribed with transmit-unconditional rates.
- Color-aware uses the incoming packet color in its algorithm
- Not promoting packets means that if the packet enters the rate limit as yellow and the rate-limit then determines that it is green, the packet remains yellow. If the rate limit determines it is red, then the packet is colored red.

A rate-limit rule is an instance of a rate-limit profile. The same profile can be used to create many rate-limit rules in the same hierarchy or in different rate-limit hierarchies. The classifier group that defines the flow can use a mark rule with color-mark profile to set the packet ToS field based on the packet color. A rate-limit hierarchy invoked from the classifier group is one way of changing the packet color; the rate-limit hierarchy is invoked before the classifier group runs the mark rule to set the packet ToS.

Hierarchical Rate-Limit Actions

Every packet traversing a rate-limit hierarchy has an owner that is defined by the last rate limit that can apply its actions to the packet; this is a configuration option.

A rate limit in the hierarchy that does not own the packet only decrements its tokens, but cannot perform any of the following actions:

- Transfer ownership of the packet to the next rate limit.
- Retain ownership of the packet but consume tokens from the remaining rate limits in the hierarchy.
- Exit the rate-limit hierarchy, making that rate limit the final one for the packet.

These actions become the same action if the hierarchy has only one rate limit. Combining these actions with the additional choices to transmit or drop packets results in the following possible actions:

- Drop—Drops the packet at that rate limit in the hierarchy. The packet does not change the state of any rate limit further down the hierarchy.
- Transmit final—Sets the packet color and ends the packet's traversal of the rate-limit hierarchy at the current rate limit. The packet is forwarded and the rate limits further down the hierarchy are not affected. Because transmit final is based on the result of the rate limit, transmit is not an attribute of the node in the rate-limit hierarchy. Committed packets can exit the hierarchy while conformed and exceeded packets continue to the next rate limit.
- Transmit conditional—Sets the packet color to the result calculated by the rate limit and forwards the packet to the next rate limit for processing, also transferring ownership of the packet to the next rate limit. The next rate limit can then set the packet color according to the state of its token buckets and apply its actions to the packet. The transmit conditional option is the same as connecting the two rate limits in series.
- Transmit unconditional—Sets the packet color to the result calculated by the rate limit, retains ownership of the packet, and forwards the packet to the next rate limit. Later rate limits only decrement their current token counts by the packet length but do not otherwise affect the packet, either by changing its color or applying their actions to it. Although the packet is not affected, the remaining rate limits change because the token counts are reduced, making them more likely to make other packets conformed or exceeded. Transmit unconditional is not allowed as an exceeded action.

After the transmit-unconditional completes, the packet traverses to the end of the hierarchy. Because ownership of the packet has been retained, no rate limit further down can apply its actions to it. Some of the later rate limits might already have very low token counts, which must still be decremented when processing a transmit-unconditional packet (if necessary, by making the token count negative). Negative token counts enable the remaining rate limits to restrict the total traffic through them to their peak rate (over a large enough averaging interval, which is a function of rates and burst sizes only). Transmit unconditional packets traversing the rate-limit hierarchy reduce the number of tokens available for other packets.

A rate limit has one of the four preceding actions configured for each possible result: committed, conformed, and exceeded. (Transmit unconditional is not allowed as an exceeded action.) The action taken depends only on the result of that rate limit, its rates, burst sizes, and current token state. In addition, the rate limit assigns a color to the packet, depending on both the result of the rate limit and the packet's incoming color. The final color after a packet has finished traversing a rate-limit hierarchy is a function of all the rate limits that owned the packet.

Policy actions are processed in the following order:

1. log
2. filter
3. traffic class
4. user packet class
5. next hop

6. rate limit
7. color status
8. color action
9. parent group
10. mark

The mark action is the last action that occurs, after parent-group, so that the color-mark profile can mark the packet with the final color from the hierarchy.



NOTE: To avoid saturation when using dual token buckets, the total amount of yellow transmit unconditional traffic should be less than the peak rate minus the committed rate; the green transmit unconditional traffic should be less than the committed rate.


Example: Multiple Flows Sharing Preferred Bandwidth Rate-Limiting Hierarchical Policy

Figure 1 on page 8 shows an interface with an attached policy that has a Video classifier that singles out a substream of the packets flowing on that interface. The Video classifier can be allocated 6 Mbps out of the 10 Mbps interface rate. All other packets on the interface are Internet. The common rate limit cannot drop Video packets, but must limit the total flow (Video and Internet) to under 10 Mbps. Internet traffic can use the Video bandwidth when there are no active Video calls, while avoiding hard partitioning of interface bandwidth.



NOTE: To avoid rate-limit saturation, we recommend that you set the rate limit profile to color-aware when the rate limit is set to receive transmit conditional.

Figure 1: Multiple Flows Sharing Preferred Bandwidth



```

host1(config)#rate-limit-profile video two-rate hierarchical
host1(config-rate-limit-profile)#committed-action transmit unconditional
host1(config-rate-limit-profile)#conformed-action transmit unconditional
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#peak-rate 60000000
host1(config-rate-limit-profile)#exit
host1(config)#rate-limit-profile common two-rate hierarchical
host1(config-rate-limit-profile)# color-aware
host1(config-rate-limit-profile)#committed-action transmit conditional
host1(config-rate-limit-profile)#conformed-action transmit conditional
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#peak-rate 100000000

```

```

host1(config-rate-limit-profile)#exit
host1(config)#policy-list mycompany
host1(config-policy-list)#classifier-group video parent-group all
host1(config-policy-list-classifier-group)#rate-limit-profile video
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#classifier-group * parent-group all
host1(config-policy-list-classifier-group)#forward
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#parent-group all
host1(config-policy-list-parent-group)#rate-limit-profile common
host1(config-policy-list-parent-group)#exit

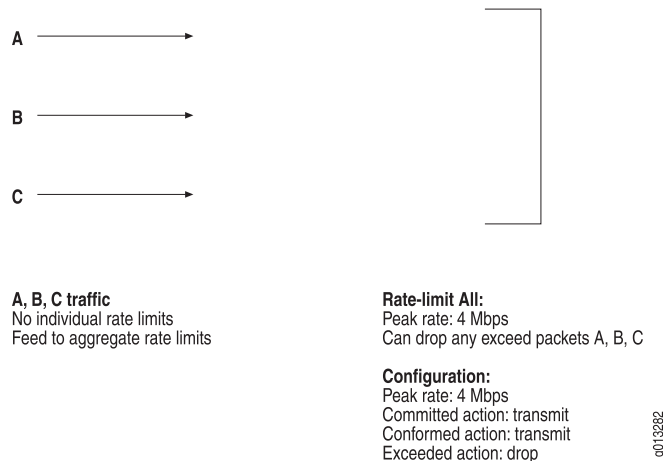
```

In this example, the rate limit Common is color-aware, using the color of the incoming packets instead of setting them to Green. This causes the rate limit Preferred to send 6 Mbps of yellow, transmit unconditional packets. The rate limit Common counts the packets against the yellow token bucket, which has a rate of 10 Mbps. However, if the rate limit Common is color-blind, it treats all packets as Green so the green token bucket gets 6 Mbps of transmit unconditional traffic, which eventually causes all packets to be saturated and dropped.

Example: Multiple Flows Sharing a Rate Limit Hierarchical Policy

Figure 2 on page 9 shows an interface that has one rate limit and three classified flows, A, B, and C. The combined traffic for A, B, and C must be below a peak rate of 4 Mbps, but each individual flow can burst up to that amount. Statistics can be collected separately on A, B, and C, while limiting only the aggregate of all three. None of the flows has any preference in accessing the rate limit, and the rate limit is shared on a first-come first-serve basis.

Figure 2: Multiple Packet Flows Sharing a Rate Limit



This example uses committed and conformed actions for a preferred rate limit profile so that the common rate limit drops only exceeded packets (those packets that raise the traffic load above 4 Mbps); packets below 4 Mbps are transmitted. By specifying **classifier-group * parent-group all**, all packets are sent to the parent group. There is no individual rate limit so that those packet use any available, unused bandwidth in the parent group rate limit.

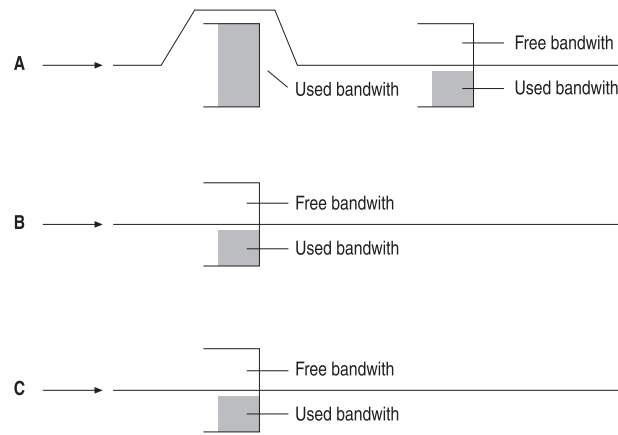
```
host1(config)#rate-limit-profile All two-rate hierarchical
host1(config-rate-limit-profile)#committed-action transmit conditional
host1(config-rate-limit-profile)#conformed-action transmit conditional
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#peak-rate 40000000
host1(config-rate-limit-profile)#exit
host1(config)#policy-list rlpshare
host1(config-policy-list)#classifier-group A parent-group All
host1(config-policy-list-classifier-group)#forward
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#classifier-group B parent-group All
host1(config-policy-list-classifier-group)#forward
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#classifier-group C parent-group All
host1(config-policy-list-classifier-group)#forward
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#parent-group All
host1(config-policy-list-parent-group)#rate-limit-profile All
host1(config-policy-list-parent-group)#exit
```

Example: Shared Pool of Additional Bandwidth with Select Flows Rate-Limiting Hierarchical Policy

Figure 3 on page 11 shows three classified flows, A, B, and C, each of which has an individual rate limit with a peak rate of 1 Mbps. If flow A is exceeding its peak rate, rather than drop the packet, the flow tries to use any bandwidth left in a shared rate limit (extrabw) of peak rate of 2 Mbps. The packet is dropped only if both the individual and the shared rate limit have no bandwidth left.

The total flow is limited to 5 Mbps, which is the sum of all the individual peak rates plus the peak rate of the shared rate limit. Individual flows A, B, and C are limited to a maximum of 3 Mbps (1 Mbps from its individual rate limit and up to 2 Mbps if it can consume the entire shared pool); however, it cannot go below a 1 Mbps rate because of the other flows. A shared rate limit enables many flows to share the extra bandwidth dynamically.

Figure 3: Shared Pool of Additional Bandwidth with Select Flows



Rate limits for A, B, C:
 Each has peak rate: 1 Mbps
 Rate limit never drops packets
 Packets under this rate transmitted with
 no further rate limiting
 Packets over this rate sent to rate-limit extrabw

Configuration:
 Peak rate: 1 Mbps
 Committed action: final
 Conformed action: final
 Exceeded action: conditional

Rate-limit extrabw:
 Each has peak rate: 2 Mbps
 Receives overflow packets from A, B, C
 Drops packets that exceed its 2 Mbps rate
 Transmits packets within 2 Mbps rate

Configuration:
 Peak rate: 2 Mbps
 Committed action: transmit
 Conformed action: transmit
 Exceeded action: drop

9013283

This example uses **transmit final** so that those packets do not pass through the common rate limit. Transmit final also indicates that there is no shared maximum. If the packets are committed or conformed, they do not need to borrow extra bandwidth or subtract tokens from it. The example uses **exceeded action transmit conditional** so that packets above the individual rate-limit maximum are not dropped but sent to the next rate limit in the hierarchy. Because this is **transmit conditional**, ownership of the packet also transfers so the common rate limit can drop these packets if it has no bandwidth left.

```
host1(config)#rate-limit-profile indiv two-rate hierarchical
host1(config-rate-limit-profile)#committed-action transmit final
host1(config-rate-limit-profile)#conformed-action transmit final
host1(config-rate-limit-profile)#exceeded-action transmit conditional
host1(config-rate-limit-profile)#peak-rate 10000000
host1(config-rate-limit-profile)#exit
host1(config)#rate-limit-profile extrabw two-rate hierarchical
host1(config-rate-limit-profile)#committed-action transmit conditional
host1(config-rate-limit-profile)#conformed-action transmit conditional
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#peak-rate 20000000
host1(config-rate-limit-profile)#exit
host1(config)#policy-list mypolicy
host1(config-policy-list)#classifier-group A parent-group extrabw
host1(config-policy-list-classifier-group)#rate-limit-profile indiv
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#classifier-group B parent-group extrabw
host1(config-policy-list-classifier-group)#rate-limit-profile indiv
host1(config-policy-list-classifier-group)#exit
```

```
host1(config-policy-list)#classifier-group C parent-group extrabw
host1(config-policy-list-classifier-group)#rate-limit-profile indiv
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#parent-group extrabw
host1(config-policy-list-parent-group)#rate-limit-profile extrabw
host1(config-policy-list-parent-group)#exit
```

Example: Aggregate Marking with Oversubscription Rate-Limiting Hierarchical Policy


Figure 4 on page 13 shows an aggregate rate limit that enables up to 2 Mbps of traffic to be sent with a ToS value marked as 10. Traffic above that rate is sent with a ToS value marked as 20 or 30 (depending on packet type) and traffic above 6 Mbps is dropped. The 2 Mbps of traffic with the ToS value of 10 is oversubscribed among individual flows A, B, and C, each of which can have up to 1 Mbps of traffic with the ToS value of 10. An individual flow can mark a packet with the ToS value of 10, but if there is insufficient bandwidth at the shared rate limit because of oversubscription, the packet is demoted and remarked.

The demoted packets from flow A are marked with the ToS value of 20 but the demoted packets from flows B and C are marked with the ToS value of 30. The shared rate limit determines whether to demote the packet, in which case each individual rate limit selects the new ToS marking. Individual flows are not required to mark demoted packets with the same value.

The committed and conformed actions are transmit conditional so that all packets also go through rate limit S, because rate limit S imposes the limit of 2 Mbps of traffic with a ToS value of 10 (total across A, B, and C).

Committed packets are transmitted conditionally to rate limit S, which has a peak rate of 6 Mbps and a committed rate of 2 Mbps; these packets can be demoted by S to Y (yellow), in which case they are remarked with the ToS value of 20 or 30. If S leaves them as G (green), they are marked with the ToS value of 10. All conformed packets from A, B, and C are also transmitted conditionally to S but arrive as Y because rate limits do not promote packets in color. S is color-aware so these Y packets do not take away G tokens, leaving them reserved only for the G packets coming from A, B, and C.

Figure 4: Aggregate Marking with Oversubscription

**Rate-limits for A, B, C:**

Packets under 1 Mbps marked TOS1

Packets between 1-2 Mbps marked with ToS value 20 (A only) or ToS value 30 (B, C)

All packets sent to rate limit S for check for ToS value 10

Rate-limit S:

Receives packets from A, B, C

Packets under 2 Mbps are not affected

Drops packets that exceed 6 Mbps rate

Demotes packets over 2 Mbps

Configuration A

Peak Rate: 2 Mbps

Committed rate: 1 Mbps

Committed action: transmit conditional

Conformed action: transmit conditional

Exceeded action: drop

G mark: TOS value 10

Y mark: **TOS value 20**R mark: **TOS value 20****Configuration B, C**

Peak Rate: 2 Mbps

Committed rate: 1 Mbps

Committed action: transmit conditional

Conformed action: transmit conditional

Exceeded action: drop

G mark: TOS value 10

Y mark: **TOS value 30**R mark: **TOS value 30****Configuration:**

Peak Rate: 6 Mbps

Committed rate: 2 Mbps

Committed action: transmit

Conformed action: transmit

Exceeded action: drop

Color-aware

9016526

```

hostl(config)#rate-limit-profile indiv two-rate hierarchical
hostl(config-rate-limit-profile)#committed-action transmit conditional
hostl(config-rate-limit-profile)#conformed-action transmit conditional
hostl(config-rate-limit-profile)#exceeded-action drop
hostl(config-rate-limit-profile)#committed-rate 10000000
hostl(config-rate-limit-profile)#peak-rate 20000000
hostl(config-rate-limit-profile)#exit
hostl(config)#rate-limit-profile S two-rate hierarchical
hostl(config-rate-limit-profile)#committed-action transmit conditional
hostl(config-rate-limit-profile)#conformed-action transmit conditional
hostl(config-rate-limit-profile)#exceeded-action drop
hostl(config-rate-limit-profile)#committed-rate 20000000
hostl(config-rate-limit-profile)#peak-rate 60000000
hostl(config-rate-limit-profile)#color-aware
hostl(config-rate-limit-profile)#exit
hostl(config)#ip color-mark-profile A
hostl(config-color-mark-profile)#green-mark 10
hostl(config-color-mark-profile)#yellow-mark 20
hostl(config-color-mark-profile)#red-mark 20
hostl(config-color-mark-profile)#exit
hostl(config)#ip color-mark-profile BC
hostl(config-color-mark-profile)# green-mark 10
hostl(config-color-mark-profile)# yellow-mark 30
hostl(config-color-mark-profile)# red-mark 30
hostl(config-color-mark-profile)# exit
hostl(config)#policy-list ToS_value_10_oversubscribed
hostl(config-policy-list)#classifier-group A parent-group S
hostl(config-policy-list-classifier-group)#rate-limit-profile indiv
hostl(config-policy-list-classifier-group)#mark profile A
hostl(config-policy-list-classifier-group)#exit
hostl(config-policy-list)#classifier-group B parent-group S
hostl(config-policy-list-classifier-group)#rate-limit-profile indiv
hostl(config-policy-list-classifier-group)#mark profile BC
hostl(config-policy-list-classifier-group)#exit

```

```
host1(config-policy-list)#classifier-group C parent-group S
host1(config-policy-list-classifier-group)#rate-limit-profile indiv
host1(config-policy-list-classifier-group)#mark profile BC
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#parent-group S
host1(config-policy-list-parent-group)#rate-limit-profile S
host1(config-policy-list-parent-group)#exit
```

Color-Aware Configuration for Rate-Limiting Hierarchical Policy

Common to many rate-limit hierarchies is a large aggregate rate limit that receives packets from many smaller individual rate limits. An individual rate limit can mark a packet yellow but, if few individual flows are active, the aggregate rate limit is likely to try to promote it to green, overriding the individual rate limit. For this reason, rate limits never promote packets in color; color-aware rate limits use the incoming color in their algorithm, but the final result is always equal to or less than the initial packet color.

Rate-limit profiles for rate-limit hierarchies include a non-default configuration option for color-aware. For two-rate rate limits this option enables the color-aware algorithm. If hierarchical, TCP-friendly one-rate rate limits have a color-aware algorithm defined.

In the following color-aware example, the non-preferred packets do not take any green tokens from rate-limit A, leaving them all for preferred packets. Preferred packets may take green and also take yellow tokens (which reduces the flow of non-preferred). In this way the non-preferred packets do not reduce the number of green preferred packets, only the number of yellow preferred packets; preferred packets are then marked from a color-mark profile.

```
class non-preferred parent A
color yellow
class preferred parent A
mark profile cm
parent A
rate-limit A !! a color-aware rate limit
```

The color-mark profile translates the packet color, which is independent of its type, to a type-dependent mark for ToS or EXP and applies it to a packet after it has exited the rate-limit hierarchy. If no translation is configured for a color, then packets of that color are not changed.

Transmit-unconditional packets entering a color-aware rate limit uses the color on the packet for the rate-limit algorithm. Doing this ensures that the color-aware rate limit depletes tokens from the token buckets to account for these packets.

Every packet sent through a rate-limit hierarchy is either dropped inside the hierarchy or emerges with a green, yellow, or red color assigned to it by the rate-limit hierarchy. The color depends on the last rate limit in the hierarchy that owned the packet and all prior rate limits. The green, yellow, or red classification applies to packets of any type and is not interface-type dependent.

A packet that has traversed the hierarchy either has been dropped or emerges with a color (green, yellow or red). This final color can be used by a mark rule with a color-mark profile to select the ToS marking for the packet. Because this operation is interface-type

dependent, the actual value is configured where the packet entered the hierarchy; however, the color is set by the entire rate-limit hierarchy.

We recommend that all rate-limit profiles that receive transmit unconditional packets should be color-aware. If not color-aware, yellow transmit unconditional packets are processed through both the green and yellow token buckets; if the green rate is low, this causes an oversubscription of transmit unconditional packets and leads to saturation. By making the rate limit color-aware, the yellow transmit unconditional packets are counted only against the yellow token bucket.

**Related
Documentation**

- [color](#)
- [color-aware](#)
- [color-mark-profile](#)
- [green-mark](#)
- [red-mark](#)
- [yellow-mark](#)

Percent-Based Rates for Rate-Limit Profiles Overview

Percent-based rate-limit profiles enable you to divide the reference rate as percentages instead of specific values. You can specify the reference rate on each interface and specify these rates in terms of percentage of this reference rate within the rate-limit profile to derive the appropriate rate. This enables you to define rate-limit profiles with rates in terms of percentage and bursts in terms of milliseconds.

You can use percent-based rate-limit profiles to:

- Configure rates in rate-limit profiles based on a percentage of a parameter. You can assign values to these parameters at the time of attachment, which enables you to use the same policy for multiple interfaces with different parameter values.
- Specify burst sizes in milliseconds when you configure percent-based rate-limits.
- Provide a generic way to configure and use policy parameters. You can use parameter names when you create policy objects and defer assigning values to these parameters until policy attachment. This enables you to share policy objects by attaching the same policy at multiple interfaces with different parameter values. You do not have to specify values each time you attach a policy; if you do not specify interface-specific, the system uses the global value.

Policy Parameter Reference-Rate

You can use a policy parameter reference-rate to derive the rates in rate-limit profiles. You can configure rate-limit profiles as a percentage of this parameter. The system calculates the rate at the time of attachment using the value assigned to this parameter for that interface.

If you do not specify a value for this parameter in Interface Configuration mode, then the Global configuration value is used.

You can modify the value of this parameter in Global Configuration mode or Interface Configuration mode. In Interface Configuration mode, you can change the value using the `increase` keyword.

If you use the `no` version of the command in Interface Configuration mode, the parameter value is set to the global default value. The `no` version of the command with the `increase` keyword decrements the value. The parameter value cannot have a negative value. The `no` version of the command in Global Configuration mode deletes the parameter if it is not used anywhere else.

Modified values affect the rates in the rate-limit profiles that are using the reference-rate parameter.



NOTE: Beginning with JunosE Release 10.3.x, you cannot modify the policy reference-rate parameter in Interface Configuration mode, if the percent-based rate-limit profile is used in external parent groups. If you attempt to change the reference-rate parameter that is referenced by multiple classification operations using the `policy-parameter reference-rate` command at the interface level, an error message is displayed in the CLI interface. This restriction exists because multiple interfaces might refer to the same external parent group resource and to prevent such interfaces from losing their reference to the older external parent group resources. However, you can modify the percent-based rate-limit profiles for external parent groups at the hierarchical rate-limit profile level.

Specifying Rates Within Rate-Limit Profiles

Within a rate-limit profile you can specify the rate either as a percentage or a specific value. In two-rate rate-limit profiles, you can select committed rate and peak rate. You can specify one rate in terms of percentage and another as a specific value. Also, one rate can be a percentage of one parameter and another rate can be a percentage of another parameter.

If the rate in a rate-limit profile is x percent, then the actual rate can be calculated from a parameter value as:

$$\text{Actual rate (in bits per second)} = (\text{parameter value} * x) / 100$$

The committed rate can be in the range 0–100 percent of the parameter value. The peak rate can be in the range 0–1000 percent of the parameter value.

The parameter value derives the appropriate rate within the rate-limit profile using a percentage. There are no validations to make the total rate less than or equal to the parameter value.

Specifying Burst Sizes

Within a rate-limit profile you can specify the burst size in milliseconds or bytes. Because rate-limit profiles have multiple rates and no restrictions, you can specify one burst in terms of milliseconds and another as bytes whether or not the corresponding rate is a percentage.

If the burst size is m milliseconds, it is calculated as:

Burst size in bytes = $(\text{rate in bps} * m) / (8 * 1000)$

In this example, the burst size can be in the range 0–10000 ms (10 seconds).

The maximum burst size is 4294967295 bytes (32 bit).

If you do not set the burst size, the system sets the default committed burst and peak burst to 100 ms. If the default burst size is less than 8192, the system changes it to 8192.

Using Service Manager with Merged Policies

When you use the Service Manager, you can attach multiple policies to the same interface point with the **merge** keyword and these policies are then merged into a new policy. The **increase** keyword enables you to change the parameter value for the profile.

If you activate the service without the **increase** keyword, the interface-specific value of the parameter is set to the value specified in the profile. However, if you activate the service with the **increase** keyword, the interface-specific value of the parameter increases by the value specified in the profile. If there was no interface-specific value at the time of activation of the profile with the **increase** keyword, then it increases from 0.

If you deactivate the service that used the **increase** keyword, the value of the parameter decreases. But if the profile did not use the **increase** keyword, deactivation does not change the current interface-specific value for that parameter. The interface-specific parameter remains until the interface is deleted.

Policy Parameter Configuration Considerations

The following list describes the rules for using policy parameters:

- Policy parameter names must be unique regardless of its type. If you configure a policy parameter with a reference-rate type, then you cannot configure it with another type until it is deleted.
- You can create policy parameters in Global Configuration mode and in Interface Configuration mode in any order.
- In Global Configuration mode, you can assign a parameter type to a parameter name and assign a default value for this parameter.
- If a parameter is configured in Global Configuration mode, but you do not assign a default value, then the system assigns a default value to the parameter. The system default value for any parameter of type reference-rate is 64K (65536).

- In Interface Configuration mode, you assign a parameter type and value for an interface. Policy parameters configured in Interface Configuration mode that have interface-type IP, IPv6, or L2TP specified with the command associate the command with the respective interface in the stack.
- If a parameter is configured in Interface Configuration mode without configuring it in Global Configuration mode, a global configuration is automatically created for this parameter with the type specified in interface configuration and a system-specified default value.
- A parameter value specified in Interface Configuration mode overrides the value specified in Global Configuration mode.
- If the parameter is not configured in Interface Configuration mode, the value from the global configuration is used. If the global value satisfies most of the interfaces, then you do not have to configure parameters for each interface separately, which reduces the number of configuration steps you need to take.
- When you delete an interface, the interface-specific configuration of the parameter is deleted. However, the global configuration remains until you delete it whether it was created explicitly in Global Configuration mode or automatically created in Interface Configuration mode.

For example, you can configure policy parameter `param1` of type `reference-rate` in Global Configuration mode with a default value of 100,000 and then configure it as 200,000 in Interface Configuration mode for `inf1`. If you configure a policy parameter as 500,000 in Interface Configuration mode for interface `inf1`, the system automatically creates parameter `param2` with a 64K (65536) global default value. When you delete interface `inf1`, the system deletes the interface-specific configuration for `param1` and `param2`, but the global configuration values of 100,000 and 64K (65536) remain until you explicitly delete them.

- You must create policy parameters in either Global Configuration mode or Interface Configuration mode before they can be used or referenced as policy objects. For example, before you define a rate in a rate-limit profile in terms of percentage of a policy parameter `param1`, you must configure `param1` as parameter type `reference-rate`.
- You can configure multiple policy parameters; there are no restrictions on the number of parameters.
- If you modify a policy parameter value in Interface Configuration mode, it affects all policies attached to that interface. If a parameter value is changed for an interface, only the input, secondary-input, and output policies attached to that interface are affected by this change.
- If you modify a policy parameter value in Global Configuration mode, it affects all policies attached to all interfaces that use the global values. For example, if parameter `param1` is used in policies attached to two interfaces, but `param1` is only configured for interface `i1`, when you modify the default value for `param1` in Global Configuration mode, it affects only the attachment on the second interface `i2`.
- You can specify a rate within a rate-limit profile as a percentage of the parameter and burst size in milliseconds. You can use this rate-limit profile in a policy. You can assign values to these parameters for an interface. The actual rate and burst size are calculated

at the time of attachment. You can attach the same policy to multiple interfaces with different parameter values.

- A hierarchical rate-limit profile that contains percentage-based parameters can be used in an external parent-group and the global default values can be changed for each external parent group instance. The following restrictions apply:
 - When you attach policies that reference such same external parent group instances to each interface, you must specify the same reference-rate policy-parameter value.
 - You cannot change the reference-rate policy-parameter values at interface level when a policy attachment to the interface exists.

One-Rate Rate-Limit Profiles Overview

E Series routers implement a single-rate rate limiter, which you can configure to provide more efficient service to TCP applications. With the single-rate rate limiter, when the committed rate is exceeded, the rate limiter drops a single packet and then resumes transmission up to a configurable burst window. The single, unacknowledged packet causes TCP to cut its transmission rate in half rather than falling back to its initial window size.



NOTE: Commands that you issue in Rate Limit Profile Configuration mode do not take effect until you exit from that mode.

The one-rate rate-limit profile attributes are:

- Color aware—Color-aware rate action (only for hierarchical rate limits)
- Committed rate—Target rate for a packet flow
- Committed burst—Amount of bandwidth allocated to accommodate bursty traffic in excess of the rate
- Excess burst—Amount of bandwidth allocated to accommodate a packet in progress when the rate is in excess of the burst
- Committed action—Drop, transmit, mark (IP and IPv6), or mark-exp (MPLS) when traffic flow does not exceed the rate; the mark value is not supported for hierarchical rate limits and the transmit values conditional, unconditional, or final are only supported on hierarchical rate limits
- Conformed action—Drop, transmit, mark (IP and IPv6), or mark-exp (MPLS) when traffic flow exceeds the rate but not the excess burst; the mark value is not supported for hierarchical rate limits and the transmit values conditional, unconditional, or final are only supported on hierarchical rate limits
- Exceeded action—Drop, transmit, mark (IP and IPv6), or mark-exp (MPLS) when traffic flow exceeds the rate; the mark value is not supported for hierarchical rate limits and the transmit values conditional, unconditional, or final are only supported on hierarchical rate limits

- Mask value—Mask to be applied with mark values for the ToS byte; applicable only to IP and IPv6 rate-limit profiles; not supported on hierarchical rate limits
- EXP mask value—Mask to be applied with mark-exp values; applicable only to MPLS rate-limit profiles; not supported on hierarchical rate limits

Two-Rate Rate-Limits Overview

The two-rate rate limiter enables you to build tiered rate-limit services and to specify different treatments for packets at different rates.

Token buckets control how many packets per second are accepted at each of the configured rates and provide flexibility in dealing with the bursty nature of data traffic. At the beginning of each sample period, the two buckets are filled with tokens based on the configured burst sizes and rates. Traffic is metered to measure its volume. When traffic is received, if tokens remain in both buckets, one token is removed from each bucket for every byte of data processed. As long as tokens are still in the committed burst bucket, the traffic is treated as committed.

When the committed burst token bucket is empty but tokens remain in the peak burst bucket, traffic is treated as conformed. When the peak burst token bucket is empty, traffic is treated as exceeded.

In color-blind mode, if the committed token bucket has enough tokens when a packet is received, the packet is green and tokens are subtracted from both the committed and the peak token buckets. If the peak bucket does not have enough tokens left, it is allowed to go negative. Green packets are the committed traffic.

If the committed bucket does not have enough tokens for the packet, the peak bucket is tested (and the committed bucket is not changed). If there are enough tokens in the peak bucket, it is decremented and the packet is yellow. Yellow packets are the conformed traffic. If the peak bucket does not have enough tokens either (because the committed bucket did not have enough tokens), the packet is red. Red packets are the exceeded traffic.

The two-rate rate-limit profile attributes are:

- ATM cell mode—ATM cell tax accounted for in statistics and rate calculations
- Color-aware—Color-aware rate action (only for hierarchical rate limits)
- Committed rate—Target rate for a packet flow
- Committed burst—Amount of bandwidth allocated to accommodate bursty traffic in excess of the committed rate
- Peak rate—Amount of bandwidth allocated to accommodate excess traffic flow over the committed rate
- Peak burst—Amount of bandwidth allocated to accommodate bursty traffic in excess of the peak rate
- Committed action—Drop, transmit, conditional, unconditional, final, mark (IP and IPv6), or mark-exp (MPLS) when traffic flow does not exceed the committed rate; the mark

value is not supported for hierarchical rate limits and the transmit values conditional, unconditional, or final are only supported on hierarchical rate limits

- Conformed action—Drop, transmit, mark (IP and IPv6), or mark-exp (MPLS) when traffic flow exceeds the committed rate but remains below the peak rate; the mark value is not supported for hierarchical rate limits and the transmit values conditional, unconditional, or final are only supported on hierarchical rate limits
- Exceeded action—Drop, transmit, mark (IP and IPv6), or mark-exp (MPLS) when traffic flow exceeds the peak rate; the mark value is not supported for hierarchical rate limits and the transmit values conditional, unconditional, or final are only supported on hierarchical rate limits
- Mask value—Mask to be applied with mark values for the ToS byte; applicable only to IP and IPv6 rate-limit profiles; not supported on hierarchical rate limits
- EXP mask value—Mask to be applied with mark-exp values; applicable only to MPLS rate-limit profiles; not supported on hierarchical rate limits

Table 3 on page 21 indicates the interaction between the rate settings and the actual traffic rate to determine the action taken by a rate-limit rule in a policy when applied to a traffic flow. This implementation is known as a two-rate, three-color marking mechanism.

Table 3: Policy Action Applied Based on Rate Settings and Traffic Rate

Peak Rate	Committed Rate = 0	Committed Rate Not 0
Peak rate = 0	<ul style="list-style-type: none"> • All traffic assigned the exceeded action 	<ul style="list-style-type: none"> • Traffic \leq committed rate assigned the committed action • Traffic $>$ committed rate assigned the exceeded action
Peak rate not 0	<ul style="list-style-type: none"> • Traffic \leq peak rate assigned the conformed action • Traffic $>$ peak rate assigned the exceeded action 	<ul style="list-style-type: none"> • Traffic \leq committed rate assigned the committed action • Committed rate $<$ Traffic $<$ peak rate assigned the conformed action • Traffic $>$ peak rate assigned the exceeded action

Table 4 on page 22 presents equations that can represent the algorithm for the two-rate rate-limit profile, where:

- B = size of packet in bytes
- Tp = size of peak token bucket in bytes (maximum size of this bucket is the configured peak burst)
- Tc = size of the committed token bucket in bytes (maximum size of this bucket is the configured committed burst)
- t = time

Table 4: Two-Rate Rate-Limit Profile Algorithms

Step	Result
If not color-aware, use green as the incoming packet color, otherwise use the actual packet color	–
If incoming packet color is green:	–
If $Tc(t) \geq B$	<ul style="list-style-type: none"> • Packet is marked as green • $Tc(t)$ is decremented by B • $Tp(t)$ is decremented by B (allow $Tp(t) < 0$ if necessary)
If $Tp(t) \geq B$ and $Tc(t) < B$	<ul style="list-style-type: none"> • Packet is marked as yellow • $Tp(t)$ is decremented by B
If $Tp(t) < B$ and $Tc(t) < B$	<ul style="list-style-type: none"> • Packet is marked as red
If incoming packet color is (only occurs in color-aware operation)	–
If $Tp(t) \geq B$	<ul style="list-style-type: none"> • Packet is marked as yellow • $Tp(t)$ is decremented by B
If $Tp(t) < B$	<ul style="list-style-type: none"> • Packet is marked as red
If incoming packet color is red (only occurs in color aware operation)	<ul style="list-style-type: none"> • Packet is marked as red

Bandwidth Management Overview

When you configure the rate-limit profile, packets are tagged with a drop preference. The color-coded tag is added automatically when the committed and peak burst values for an interface's rate-limit profile are exceeded. The egress forwarding controller uses the drop preference to determine which packets are dropped when there is contention for outbound queuing resources within the E Series router.

The queuing system uses drop eligibility to select packets for dropping when congestion exists on an egress interface. This method is called dynamic color-based threshold dropping. The 2-bit tag assigns a color code to the packet: red, yellow, or green. Each packet queue has two color-based thresholds as well as a queue limit:

- Red packets are dropped when congestion causes the queue to fill above the red threshold.
- Yellow packets are dropped when the yellow threshold is reached.
- Green packets are dropped when the queue limit is reached.

This internal tagging is done automatically when a rate-limit profile is applied to an interface and does not necessarily reflect the operation of the policy on an interface.

Having a committed rate and a peak rate enables you to configure two different fill rates for the token buckets. For example, you can configure the fill rate on the peak token bucket to be faster than the fill rate on the committed bucket. This configuration enables you to accommodate bursts of traffic, but, through coloring, it enables you to identify which packets are committed and which ones are not.

To enforce ingress data rates below the physical line rate of a port, you can rate limit a classified packet flow at ingress. A rate-limit profile with a policy rate-limit profile rule provides this capability. The rate-limit profile defines the attributes of the desired rate.


You can set an action based on one rate or two rates. These actions include drop, transmit, or mark. The default is to transmit committed and conformed packets, and to drop exceeded packets.

A color-coded tag is added automatically to each packet based on the following categories:

- Committed—Green
- Conformed—Yellow
- Exceeded—Red

Figure 5 on page 23 illustrates congestion management.

Figure 5: Congestion Management



Examples: One-Rate Rate-Limit Profile

A one-rate rate-limit profile can be configured for hard tail drop rate-limit or TCP-friendly behavior. Packets can be categorized as committed, conformed, or exceeded.

You can configure a one-rate rate-limit profile to hard limit a packet flow to a specified rate. To rate limit the traffic on an interface from source IP address 1.1.1.1 to 1 Mbps, issue the following commands:

```
host1#configure terminal
host1(config)#ip rate-limit-profile oneMegRlp one-rate
host1(config-rate-limit-profile)#committed-rate 1000000
host1(config-rate-limit-profile)#exit
host1(config)#ip classifier-list claclA ip host 1.1.1.1 any
host1(config)#ip policy-list testPolicy
host1(config-policy-list)#classifier-group claclA
host1(config-policy-list-classifier-group)#rate-limit-profile oneMegRlp
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#exit
host1(config)#interface atm 0/0.0
host1(config-subif)#ip policy input testPolicy statistics enabled
```

You can also configure a one-rate rate-limit profile to provide a TCP-friendly rate limiter. To configure a rate limiter with TCP-friendly characteristics, we recommend that you set the committed burst to allow for 1 second of data at the specified rate, and the excess burst to allow 1.5 seconds of data at the specified committed rate plus the committed burst. For example:

```
host1(config)#ip rate-limit-profile tcpFriendly8MB one-rate
host1(config-rate-limit-profile)#committed-rate 8000000
host1(config-rate-limit-profile)#committed-burst 1000000
host1(config-rate-limit-profile)#excess-burst 2500000
host1(config-rate-limit-profile)#committed-action transmit
host1(config-rate-limit-profile)#exceeded-action drop
```

Examples: Two-Rate Rate-Limit Profile

You can configure a two-rate rate-limit profile for two different rates, committed and peak, that are used to define a two-rate, three-color marking mechanism. You can categorize packets as committed, conformed, or exceeded:

- Up to the committed rate, packets are considered to be committed.
- From the committed to peak rate, packets are considered to be conformed.
- After the peak rate, packets are considered to be exceeded.

This configuration is implemented with token buckets. See RFC 2698 for more details.

The following example rate limits traffic on an interface from source IP address 1.1.1.1 so that traffic at a rate up to 1 Mbps is colored green and transmitted, traffic at a rate from 1 Mbps to 2 Mbps is colored yellow and transmitted, and traffic at a rate above 2 Mbps is dropped.

```
host1(config)#ip rate-limit-profile 1MbRLP
host1(config-rate-limit-profile)#committed-rate 1000000
host1(config-rate-limit-profile)#peak-rate 2000000
host1(config-rate-limit-profile)# committed-action transmit
host1(config-rate-limit-profile)#conformed-action transmit
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#exit
host1(config)#ip classifier-list claclA ip host 1.1.1.1 any
host1(config)#ip policy-list testPolicy
host1(config-policy-list)#classifier-group claclA
host1(config-policy-list-classifier-group)#rate-limit-profile 1MbRLP
host1(config-policy-list-classifier-group)# exit
host1(config-policy-list)#exit
host1(config-policy-list)#interface atm 0/0.0
host1(config-subif)#ip policy input testPolicy statistics enabled
```

Examples: Rate-Limiting Individual or Aggregate Packet Flows

You can construct policies to provide rate limiting for individual packet flows or for the aggregate of multiple packet flows. For example, if you have traffic from multiple sources, you can either rate limit each traffic flow individually, or you can rate limit the aggregate flow for the traffic from all sources.

- To rate limit individual packet flows, use a separate classifier list to classify each flow.
- To rate limit the aggregate of multiple traffic flows, use a single classifier list for the multiple entries.

In the following example, interface ATM 3/1.1 classifies on three traffic flows from different sources. Each traffic flow is rate limited to 1MB (which is defined by the rate-limit profile rl1Meg).

```
host1(config)#ip classifier-list clFlow1 ip host 10.1.1.1 any
host1(config)#ip classifier-list clFlow2 ip host 10.1.1.2 any
host1(config)#ip classifier-list clFlow3 ip host 10.1.1.3 any
host1(config)#ip policy-list plRateLimit
host1(config-policy-list)#classifier-group clFlow1
host1(config-policy-list-classifier-group)#rate-limit-profile rl1Meg
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#classifier-group clFlow2
host1(config-policy-list-classifier-group)#rate-limit-profile rl1Meg
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#classifier-group clFlow3
host1(config-policy-list-classifier-group)#rate-limit-profile rl1Meg
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#exit
host1(config)#interface atm 3/1.1
host1(config-subif)#ip policy input plRateLimit statistics enabled
host1(config-subif)#exit
```

In the following example, interface ATM 3/1.1 again classifies on three traffic flows; however, this policy rate limits the aggregate of the three flows to 1 MB.

```
host1(config)#ip classifier-list clFlowAll ip host 10.1.1.1 any
host1(config)#ip classifier-list clFlowAll ip host 10.1.1.2 any
host1(config)#ip classifier-list clFlowAll ip host 10.1.1.3 any
host1(config)#ip policy-list plRateLimit
host1(config-policy-list)#classifier-group clFlowAll
host1(config-policy-list-classifier-group)#exit
host1(config-policy-list)#exit
host1(config)#interface atm 3/1.1
host1(config-subif)#ip policy input plRateLimit statistics enabled
host1(config-subif)#exit
```


PART 2

Configuration

- [Configuration Tasks for Rate-Limit Profile Management on page 29](#)
- [Configuration Tasks for One-Rate Rate-Limit Profile Management on page 39](#)
- [Configuration Tasks for Two-Rate Rate-Limit Profile Management on page 45](#)

CHAPTER 2

Configuration Tasks for Rate-Limit Profile Management

- [Creating Rate-Limit Profiles on page 29](#)
- [Policy Parameter Quick Configuration on page 34](#)
- [Setting the Committed Action for a Rate-Limit Profile on page 34](#)
- [Setting the Committed Burst for a Rate-Limit Profile on page 35](#)
- [Setting the Committed Rate for a Rate-Limit Profile on page 36](#)
- [Setting the Conformed Action for a Rate-Limit Profile on page 36](#)
- [Setting the Exceeded Action for a Rate-Limit Profile on page 36](#)
- [Setting the Excess Burst for a Rate-Limit Profile on page 37](#)
- [Setting the Mask Value for MPLS Rate-Limit Profiles on page 37](#)
- [Setting the Mask Value for IP and IPv6 Rate-Limit Profiles on page 37](#)
- [Setting the Peak Rate for Rate-Limit Profiles on page 37](#)
- [Rate-Limiting Traffic Flows on page 38](#)

Creating Rate-Limit Profiles

Create rate-limit profiles with a rate based on percentage and a burst in milliseconds. The system creates a policy using these rate-limit profiles and then attaches them to different interfaces using different parameter values.

1. Create policy parameter refRlpRate.

```
host1(config)#policy-parameter refRlpRate reference-rate
host1(config-policy-param-reference-rate)#reference-rate 100000
host1(config-policy-param-reference-rate)#exit
```

2. Create rate-limit profile rlpData.

```
host1(config)#ip rate-limit-profile rlpData
host1(config-rate-limit-profile)#committed-rate refRlpRate percentage 10
host1(config-rate-limit-profile)#committed-burst millisecond 100
host1(config-rate-limit-profile)#peak-rate refRlpRate percentage 100
host1(config-rate-limit-profile)#peak-burst millisecond 150
host1(config-rate-limit-profile)#exit
```

3. Create rate-limit profile rlpVoice.

```

host1(config)#ip rate-limit-profile rlpVoice
host1(config-rate-limit-profile)#committed-rate 64000
host1(config-rate-limit-profile)#committed-burst 100000
host1(config-rate-limit-profile)#peak-rate refRlpRate percentage 100
host1(config-rate-limit-profile)#peak-burst millisecond 150
host1(config-rate-limit-profile)#exit

```

4. Create rate-limit profile rlpVideo.

```

host1(config)#ip rate-limit-profile rlpVideo
host1(config-rate-limit-profile)#committed-rate refRlpRate percentage 70
host1(config-rate-limit-profile)#committed-burst millisecond 100
host1(config-rate-limit-profile)#peak-rate refRlpRate percentage 100
host1(config-rate-limit-profile)#peak-burst millisecond 150
host1(config-rate-limit-profile)#exit

```

5. Create the policy.

```

host1(config)#ip policy-list P
host1(config-policy)#classifier-group data
host1(config-policy-classifier-group)#rate-limit-profile rlpData
host1(config-policy-classifier-group)#exit
host1(config-policy)#classifier-group voice
host1(config-policy-classifier-group)#rate-limit-profile rlpVoice
host1(config-policy-classifier-group)#exit
host1(config-policy)#classifier-group video
host1(config-policy-classifier-group)#rate-limit-profile rlpVideo
host1(config-policy-classifier-group)#exit
host1(config-policy)#exit

```

6. Attach IP Policy P at interface atm5/0.1.

```

host1(config)#interface atm 5/0.1
host1(config-if)#ip policy-parameter reference-rate refRlpRate 1000000
host1(config-if)#ip policy input P

```

7. Attach IP Policy P at interface atm5/0.2 with merge.

```

host1(config)#interface atm 5/0.2
host1(config-if)#ip policy input P stats enabled merge

```

8. Display the policy list.

```

host1#show policy-list

```

```

Policy Table
-----

```

```

IP Policy P

```

```

Administrative state: enable
Reference count:      1
Classifier control list: data, precedence 100
    rate-limit-profile rlpData
Classifier control list: voice, precedence 100
    rate-limit-profile rlpVoice
Classifier control list: video, precedence 100
    rate-limit-profile rlpVideo

```

```

Referenced by interfaces:

```

```

    ATM5/0.1 input policy, statistics disabled, virtual-router default
    ATM5/0.2 input policy, statistics enabled, virtual-router default

```

Referenced by profiles:
None

Referenced by merge policies:
None

9. Display the rate-limit profiles.

host1#show rate-limit-profile

Rate Limit Profile Table

IP Rate-Limit-Profile: rlpData

Profile Type: two-rate
Reference count: 1
Committed rate: refRlpRate % 10
Committed burst: 100 milliseconds
Peak rate: refRlpRate % 100
Peak burst: 150 milliseconds
Mask: 255
Committed rate action: transmit
Conformed rate action: transmit
Exceeded rate action: drop

IP Rate-Limit-Profile: rlpVoice

Profile Type: two-rate
Reference count: 1
Committed rate: 64000
Committed burst: 100000
Peak rate: refRlpRate % 100
Peak burst: 150 milliseconds
Mask: 255
Committed rate action: transmit
Conformed rate action: transmit
Exceeded rate action: drop

IP Rate-Limit-Profile: rlpVideo

Profile Type: two-rate
Reference count: 1
Committed rate: refRlpRate % 70
Committed burst: 100 milliseconds
Peak rate: refRlpRate % 100
Peak burst: 150 milliseconds
Mask: 255
Committed rate action: transmit
Conformed rate action: transmit
Exceeded rate action: drop

10. Display policy parameters. If a rate-limit profile uses this parameter twice then it increases the reference count by 2.

host1#show policy-parameter brief

Reference-rate refRlpRate: 100000, 6 references

Display policy parameters

host1#show policy-parameter

Policy Parameter refRlpRate

Type: reference-rate
Rate: 100000
Reference count: 6
Referenced by interfaces: 1 references
IP interface ATM5/0.1: 1000000

```
Referenced by rate-limit profiles: 5 references
rlpData
rlpVoice
rlpVideo
```

11. Display interface atm5/0.1.

```
host1#show ip interface atm 5/0.1
ATM5/0.1 line protocol Atm1483 is down, ip is down (ready)
Network Protocols: IP
Internet address is 1.1.1.1/255.255.255.255
Broadcast address is 255.255.255.255
Operational MTU = 0 Administrative MTU = 0
Operational speed = 1000000000 Administrative speed = 0
Discontinuity Time = 0
Router advertisement = disabled
Proxy ARP = disabled
Network Address Translation is disabled
TCP MSS Adjustment = disabled
Administrative debounce-time = disabled
Operational debounce-time = disabled
Access routing = disabled
Multipath mode = hashed
Auto Configure = disabled
Auto Detect = disabled
Inactivity Timer = disabled

In Received Packets 0, Bytes 0
  Unicast Packets 0, Bytes 0
  Multicast Packets 0, Bytes 0
In Policed Packets 0, Bytes 0
In Error Packets 0
In Invalid Source Address Packets 0
In Discarded Packets 0
Out Forwarded Packets 0, Bytes 0
  Unicast Packets 0, Bytes 0
  Multicast Routed Packets 0, Bytes 0
Out Scheduler Dropped Packets 0, Bytes 0
Out Policed Packets 0, Bytes 0
Out Discarded Packets 0

IP policy input P
  Statistics are disabled
```

12. Display interface atm5/0.2.

```
host1#show ip interface atm 5/0.2
ATM5/0.2 line protocol Atm1483 is down, ip is down (ready)
Network Protocols: IP
Internet address is 2.2.2.2/255.255.255.255
Broadcast address is 255.255.255.255
Operational MTU = 0 Administrative MTU = 0
Operational speed = 1000000000 Administrative speed = 0
Discontinuity Time = 0
Router advertisement = disabled
Proxy ARP = disabled
Network Address Translation is disabled
TCP MSS Adjustment = disabled
Administrative debounce-time = disabled
Operational debounce-time = disabled
Access routing = disabled
Multipath mode = hashed
```

```

Auto Configure = disabled
Auto Detect = disabled
Inactivity Timer = disabled

In Received Packets 0, Bytes 0
  Unicast Packets 0, Bytes 0
  Multicast Packets 0, Bytes 0
In Policed Packets 0, Bytes 0
In Error Packets 0
In Invalid Source Address Packets 0
In Discarded Packets 0
Out Forwarded Packets 0, Bytes 0
  Unicast Packets 0, Bytes 0
  Multicast Routed Packets 0, Bytes 0
Out Scheduler Dropped Packets 0, Bytes 0
Out Policed Packets 0, Bytes 0
Out Discarded Packets 0

IP policy input P
  classifier-group data entry 1
    0 packets, 0 bytes
    rate-limit-profile rlpData
      committed rate: 10000 bps, committed burst: 125 bytes
      peak Rate: 100000 bps, peak burst: 1875 bytes
      committed: 0 packets, 0 bytes, action: transmit
      conformed: 0 packets, 0 bytes, action: transmit
      exceeded: 0 packets, 0 bytes, action: drop
  classifier-group voice entry 1
    0 packets, 0 bytes
    rate-limit-profile rlpVoice
      committed rate: 64000 bps, committed burst: 100000 bytes
      peak Rate: 100000 bps, peak burst: 1875 bytes
      committed: 0 packets, 0 bytes, action: transmit
      conformed: 0 packets, 0 bytes, action: transmit
      exceeded: 0 packets, 0 bytes, action: drop
  classifier-group video entry 1
    0 packets, 0 bytes
    rate-limit-profile rlpVideo
      committed rate: 70000 bps, committed burst: 875 bytes
      peak Rate: 100000 bps, peak burst: 1875 bytes
      committed: 0 packets, 0 bytes, action: transmit
      conformed: 0 packets, 0 bytes, action: transmit
      exceeded: 0 packets, 0 bytes, action: drop

```

To configure a policy-parameter at an interface with the **increase** keyword:

1. Create policy list P2.

```

host1(config)#ip policy-list P2
host1(config-policy)#classifier-group data2
host1(config-policy-classifier-group)#rate-limit-profile rlpData
host1(config-policy-classifier-group)#exit
host1(config-policy)#exit

```

2. Attach IP Policy P2 at interface atm5/0.2 with the **merge** keyword.

```

host1(config)#interface atm 5/0.2
host1(config-if)#ip policy-parameter reference-rate refRlpRate 100000

```

This increases from 0.

```
host1(config)#ip policy-parameter reference-rate refRlpRate increase 100000
```

This increases from the existing 100000.

```
host1(config)#ip policy input P2 merge
```

3. Verify the configuration.

```
host1#show policy-parameter
Policy Parameter refRlpRate
  Type: reference-rate
  Rate: 100000
  Reference count: 7
  Referenced by interfaces: 2 references
    IP interface ATM5/0.1: 1000000
    IP interface ATM5/0.2: 200000
  Referenced by rate-limit profiles: 5 references
    rlpData
    rlpVoice
    rlpVideo
```

Policy Parameter Quick Configuration

To configure policing, use the following steps:

1. Configure a policy parameter in Global Configuration mode.
2. Assign the parameter type and global default value to a parameter.
3. Use this policy parameter in policy objects, create a generic policy, and attach it to multiple interfaces.
4. Adjust the policy parameter value for a specific interface by configuring it in Interface Configuration mode for any interface.

Setting the Committed Action for a Rate-Limit Profile

You can use the **committed-action** command to set the committed action for a rate-limit profile. Packets are colored green. For IP and IPv6 rate-limit profiles, mark the packet by setting the ToS byte (IP) or traffic class field (IPv6) to the specified 8-bit value, and transmit the packet. The mark value is masked with the default 255 unless it is overridden by the **mask-val** command to specify a different mask; not supported on hierarchical rate limits. For MPLS rate-limit profiles, set the EXP bits of MPLS packets to the specified value in the range 0–7, and transmit the packet. The mark EXP value is masked with the default 7 unless you use the **exp-mask** command to specify a different mask; not supported on hierarchical rate limits. The **no** version restores the default value, **transmit**



NOTE: If a rate-limit rule and a mark or mark-exp rule have the same classifier, the marking set by a rate-limit rule takes precedence when packets are processed by that classifier. The rate-limit rule marking overwrites any marking set by the mark rule (to the ToS byte) or by the mark-exp rule (to the EXP bits). These actions do not apply to hierarchical rate limits.

To configure the committed action, enter Rate Limit Profile Configuration mode.

- Issue the **committed-action** command:

```
host1(config-rate-limit-profile)#committed-action transmit
```

**Related
Documentation**

- committed-action

Setting the Committed Burst for a Rate-Limit Profile

You can use the **committed-burst** command to set the committed burst in bytes; range is 1–4294967295. You can use the **committed-burst** command to set the committed burst in milliseconds for a rate-limit profile; range is 1–10000. The no version restores the default value, 8192 bytes if the rate is in bytes per second; 100 milliseconds if the rate is in milliseconds.

When you specify a nonzero value for the rate, the burst size is automatically calculated for a 100-ms burst as described for the **committed-rate** command. If the calculated burst size is less than the default value of 8 KB, the default value (8192 bytes) is used.



NOTE: We recommend that you do not configure a committed or peak burst size smaller than the MTU of the interface. Doing so causes large packets to be dropped even when they are transmitted at a very low rate.

When you specify a nonzero value for the committed rate, the committed burst size is calculated based on a 100-ms burst as follows:

committed burst in bytes = (committed rate in bps x 100 ms) ÷ 8 bits per byte

The router displays committed rate in bits per second and committed burst in bytes. For example, if the rate is 8 Mbps, the burst size is 100 ms x 8 Mbps = 800,000 bits or 100,000 bytes:

committed burst = (8,000,000 bps x 100 ms) ÷ 8 = 100,000 bytes

For this example, displaying the rate-limit profile shows:

```
committed-rate 8000000
```

```
committed-burst 100000
```

If the calculated burst value is less than the default burst size of 8 KB, the default burst size is used. For most configurations this value probably is sufficient, making it optional for you to configure a value for the associated committed burst size.

To configure the committed burst, enter Rate Limit Profile Configuration mode.

- Issue the **committed-burst** command:

```
host1(config-rate-limit-profile)#committed-burst 20000
```

Related Documentation

- committed-burst

Setting the Committed Rate for a Rate-Limit Profile

You can set the committed rate as a percentage of a reference rate defined in the specified policy parameter.

- Issue the **committed-rate** command from Rate Limit Profile Configuration mode to set the committed rate in bits per second for a rate-limit profile:

```
host1(config-rate-limit-profile)#committed-rate refRlpRate percentage 10
```

Related Documentation

- committed-rate

Setting the Conformed Action for a Rate-Limit Profile

You can use the **conformed-action** command. Packets are colored yellow. For IP and IPv6 rate-limit profiles, mark the packet by setting the ToS byte (IP) or traffic class field (IPv6) to the specified 8-bit value, and transmit the packet. The mark value is masked with the default 255 unless it is overridden by the **mask-val** command to specify a different mask; not supported on hierarchical rate limits. For MPLS rate-limit profiles, set the EXP bits of MPLS packets to the specified value in the range 0–7, and transmit the packet. The mark EXP value is masked with the default 7 unless you use the **exp-mask** command to specify a different mask; not supported on hierarchical rate limits. To set the conformed action for a rate-limit profile:

- Issue the **conformed-action** command from Rate Limit Profile Configuration mode:

```
host1(config-rate-limit-profile)#conformed-action transmit
```

Related Documentation

- conformed-action

Setting the Exceeded Action for a Rate-Limit Profile

You can use the **exceeded-action** command to set the exceeded action for a rate-limit profile: Packets are colored red. For IP and IPv6 rate-limit profiles, mark the packet by setting the ToS byte (IP) or traffic class field (IPv6) to the specified 8-bit value, and transmit the packet. The mark value is masked with the default 255 unless it is overridden by the **mask-val** command to specify a different mask; not supported on hierarchical rate limits. For MPLS rate-limit profiles, set the EXP bits of MPLS packets to the specified value in the range 0–7, and transmit the packet. The mark EXP value is masked with the default 7 unless you use the **exp-mask** command to specify a different mask; not supported on hierarchical rate limits. The **no** version restores the default value, **drop**.

- Issue the **exceeded-action** command from Rate Limit Profile Configuration mode:

```
host1(config-rate-limit-profile)#exceeded-action drop
```


Related Documentation

- [exceeded-action](#)

Setting the Excess Burst for a Rate-Limit Profile

For one-rate rate-limit profiles only, use the **excess-burst** command to set the excess burst in bytes for a rate-limit profile; range is 0–4294967295. Use the **excess-burst** command to set the excess burst in milliseconds for a rate-limit profile; range is 1–10000. The **no** version restores the default value, 0.

- Issue the **excess-burst** command from Rate Limit Profile Configuration mode:

```
host1(config-rate-limit-profile)#excess-burst millisecond 1000
```

Related Documentation

- [excess-burst](#)

Setting the Mask Value for MPLS Rate-Limit Profiles

You can use the **exp-mask** command to set the mask value used for MPLS rate-limit profiles, in the range 1–255. The **no** version restores the default value, 7. This command is associated with the **committed-action**, **conformed-action**, and **exceeded-action** commands.

- Issue the **exp-mask** command from Rate Limit Profile Configuration mode.

```
host1(config-rate-limit-profile)#exp-mask 5
```

Related Documentation

- [mask-val](#)

Setting the Mask Value for IP and IPv6 Rate-Limit Profiles

You can use the **mask-val** command to set the mask value used for IP and IPv6 rate-limit profiles. 5.2.0b1 ID-2170 Use the mask values to set the appropriate bits in the ToS field of the IP packet header or in the traffic class field of the IPv6 packet header. The **no** version restores the default value, 255. This command is associated with the **committed-action**, **conformed-action**, and **exceeded-action** commands.

- Issue the **mask-val** command from Rate Limit Profile Configuration mode:

```
host1(config-rate-limit-profile)#mask-val 0xFC
```

Related Documentation

- [mask-val](#)

Setting the Peak Rate for Rate-Limit Profiles

For two-rate rate-limit profiles only, you can use the **peak-rate** command to set the peak rate in bits per second for a rate-limit profile; range is 1–4294967295. Use to set the peak rate as a percentage value; range is 0–100. During a software upgrade, the peak

rate in a rate-limit profile is automatically set to 0 if it was nonzero but less than the committed rate before the upgrade. The **no** version restores the default value, 0.

When you specify a nonzero value for the peak rate, the peak burst size is calculated based on a 100-ms burst as follows:

$$\text{peak burst in bytes} = (\text{peak rate in bps} \times 100 \text{ ms}) \div 8 \text{ bits per byte}$$

The CLI displays peak rate in bits per second and peak burst in bytes. For example, if the rate is 8 Mbps, the burst size is $100 \text{ ms} \times 8 \text{ Mbps} = 800,000 \text{ bits}$ or 100,000 bytes:

$$\text{peak burst} = (8,000,000 \text{ bps} \times 100 \text{ ms}) \div 8 = 100,000 \text{ bytes}$$

For this example, displaying the rate-limit profile shows:

```
peak-rate 8000000
peak-burst 100000
```

If the calculated peak burst value is less than the default peak burst size of 8 KB, the default burst size is used. For most configurations this value is probably sufficient, making it optional to configure the associated peak burst size.

- Issue the **peak-rate** command in Rate Limit Profile Configuration mode to set the peak rate:

```
host1(config-rate-limit-profile)#peak-rate refRlpRate percentage 100
```

Related Documentation

- [peak-rate](#)

Rate-Limiting Traffic Flows

You can rate limit traffic flows destined for an SRP module by implementing a token bucket policer. The configured rate limits are stored in NVS and persist across reboots.

Related Documentation

- [Rate Limits for Interfaces Overview on page 3](#)
- [Monitoring Policy Management Overview](#)
- [control-plane](#)
- [policer](#)

CHAPTER 3

Configuration Tasks for One-Rate Rate-Limit Profile Management

- [Creating a One-Rate Rate-Limit Profile on page 39](#)
- [Configuring a TCP-Friendly One-Rate Rate-Limit Profile on page 40](#)
- [Setting a One-Rate Rate-Limit Profile on page 42](#)

Creating a One-Rate Rate-Limit Profile

To create or modify a one-rate rate-limit profile, use the following commands with the **one-rate** keyword:

- **ip** rate-limit-profile command
- **ipv6** rate-limit-profile command
- **l2tp** rate-limit-profile command
- **mpls** rate-limit-profile command

The following example creates a rate-limit profile named `tcpFriendly8Mb`. This rate-limit profile, when included as part of a rule in a policy list, sets a TCP-friendly rate for a specified flow:

```
host1(config)#ip rate-limit-profile tcpFriendly8Mb one-rate
host1(config-rate-limit-profile)#committed-rate 8000000
host1(config-rate-limit-profile)#committed-burst 1500000
host1(config-rate-limit-profile)#excess-burst 3000000
host1(config-rate-limit-profile)#committed-action transmit
host1(config-rate-limit-profile)#conformed-action transmit
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#mask-val 255
```

To configure a single-rate hard limit, set the committed rate and burst rate to the desired values, the committed action to transmit, the conformed action to drop, and the exceeded action to drop. The peak rate must be set to zero.



NOTE: You can also achieve the characteristics of the single-rate hard limit by configuring a one-rate rate-limit profile with the extended burst rate set to zero.

- Related Documentation**
- [Rate Limits for Interfaces Overview on page 3](#)
 - [Monitoring Policy Management Overview](#)

Configuring a TCP-Friendly One-Rate Rate-Limit Profile

You can configure a committed rate, committed burst, and excess burst for the token bucket. For example, to configure a rate-limit process with hard tail dropping of packets when tokens are unavailable, set the committed rate and committed burst to a nonzero value, and set the excess burst to zero. Setting the excess burst to a nonzero value causes the router to drop packets in a more friendly way.

The configuration values for the preceding attributes determine the degree of friendliness of the rate-limit process. Instead of tail dropping packets that arrive outside the committed and burst rate envelope, the TCP-friendly bucket enables more tokens to be borrowed, up to a limit determined by the excess burst size. The next packet that borrows tokens in excess of the excess burst size is deemed excessive and is dropped if the exceeded action is set to drop.

The rate-limit algorithm is designed to avoid consecutive packet drops in the initial stages of congestion when the packet flow rate exceeds the committed rate of the token bucket. The intention is that just a few packet drops are sufficient for TCP's congestion control algorithm to drastically scale back its sending rate. Eventually, the packet flow rate falls below the committed rate, which enables the token bucket to replenish faster because of the reduced load.

If the packet flow rate exceeds the committed rate for an extended period of time, the rate-limit algorithm tends toward hard tail dropping. In a properly configured scenario, the rate limiter is consistently driven to borrow tokens because of TCP's aggressive nature, but it replenishes the tokens as TCP backs off, resulting in a delivered rate that is very close to the rate configured in the rate-limit profile.

The recommended burst sizes for TCP-friendly behavior are:

- Committed burst—0.2 to 2.0 seconds of the committed rate
- Excess burst—1.0 to 2.0 seconds of the committed rate, plus the committed burst

For example, if the committed rate is 1,000,000 bps, the recommended burst sizes are as follows:

- Committed burst is $1,000,000 \times 1.0 \times 1/8 = 125,000$ bytes

Multiplying the committed rate by 1.0 seconds converts the rate to bits, then multiplying the number of bits by 1/8 converts the value to bytes.

- Excess burst is $1,000,000 \times 1.5 \times 1/8 + 125,000 = 312,500$ bytes

Multiplying the committed rate by 1.5 converts the rate to bits, then multiplying the number of bits by 1/8 converts the value to bytes.

TCP-friendly rate limits have only one token bucket, but they also maintain a cumulative debt counter that represents how much traffic above the committed rate has recently been seen. This cumulative debt increases until it reaches the extended burst value; at that point the cumulative debt is reset to 0, but the offending packet is marked red. The cumulative debt increases faster than just by the packet size, so if the TCP source does not respond to TCP flow control and more of its packets are dropped.

[Table 5 on page 41](#) presents equations that can also represent the algorithm for the TCP-friendly one-rate rate limit profile when using hierarchical rate limiting, where:

- B = size of packet in bytes
- CD = cumulative debt
- t = time
- $T(t)$ = number of tokens in token bucket at time t

Table 5: TCP-Friendly One-Rate Rate-Limit Profile Algorithms

Step	Result
If not color aware, use green as the incoming packet color, otherwise use the actual packet color	—
If incoming packet color is green	—
If $T(t) \geq B$	<ul style="list-style-type: none"> • Packet is colored • $T(t)$ is decremented by B
If $T(t) < B$ and CD is incremented by $B - T(t)$	—
If $CD < \text{Extended Burst}$ and $T(t) < B$	<ul style="list-style-type: none"> • Packet is colored • $T(t)$ is decremented by B (allow $T(t) < 0$, if necessary)
If $CD \geq \text{Extended Burst}$ and $T(t) < B$	<ul style="list-style-type: none"> • Packet is colored • CD is reset to 0
If incoming packet color is (only occurs in color-aware operation)	—
If $T(t) < B$ and CD is incremented by $B - T(t)$	—
If $CD < \text{Extended Burst}$	<ul style="list-style-type: none"> • Packet is colored yellow • $T(t)$ is decremented by B (allow $T(t) < 0$, if necessary)
If $CD \geq \text{Extended Burst}$	<ul style="list-style-type: none"> • Packet is colored • CD is reset to 0
If incoming packet color is (only occurs in color-aware operation)	<ul style="list-style-type: none"> • Packet is colored red

Setting a One-Rate Rate-Limit Profile

You can use the **rate-limit-profile one-rate** command to create a rate-limit profile and enter Rate Limit Profile Configuration mode, from which you can configure attributes for the rate-limit profile. See [Table 3 on page 21](#).



NOTE: The JunosE Software includes the layer 2 headers in the calculations it uses to enforce the rates that you specify in rate-limit profiles.

Use one of the **ip**, **ipv6**, **l2tp**, or **mpls** keywords in front of the command to specify the type of rate-limit profile you want to create or modify. If you do not include one of the keywords, the router creates an IP rate-limit profile by default.

For hierarchical rate limits, do not specify the interface type, but add the **hierarchical** keyword at the end. The **color-aware** keyword is only supported on hierarchical rate limits.

If you do not include a **one-rate** or **two-rate** keyword, the default is a two-rate rate-limit profile. If you enter a **rate-limit-profile** command with the **one-rate** keyword and then type **exit**, the router creates a rate-limit profile with the default values listed in [Table 6 on page 42](#).

Table 6: One-Rate Rate-Limit-Profile Defaults

Policy Attribute	Default Value
type	one-rate
committed-rate	0
committed-burst	8192
excess-burst	0
committed-action	transmit
conformed-action	transmit
exceeded-action	drop
mask (IP and IPv6 rate-limit profiles)	255
exp-mask (MPLS rate-limit profiles)	7



NOTE: We recommend that you do not configure a committed or peak burst size smaller than the MTU of the interface. Doing so causes large packets to be dropped even when they are transmitted at a very low rate.

- Issue the **ip rate-limit-profile** command in Global Configuration mode:

```
host1(config)#ip rate-limit-profile tcpFriendly10Mb one-rate
```



NOTE: Commands that you issue in Rate Limit Profile Configuration mode do not take effect until you exit from that mode.

**Related
Documentation**

- [rate-limit-profile](#)

CHAPTER 4

Configuration Tasks for Two-Rate Rate-Limit Profile Management

- [Creating a Two-Rate Rate-Limit Profile on page 45](#)
- [Setting the Peak Burst for Two-Rate Rate-Limit Profiles on page 46](#)
- [Setting a Two-Rate Rate-Limit-Profile on page 46](#)

Creating a Two-Rate Rate-Limit Profile

To create or modify a two-rate rate-limit profile, use the following commands with the **two-rate** keyword:

- **ip** rate-limit-profile command
- **ipv6** rate-limit-profile command
- **l2tp** rate-limit-profile command
- **mpls** rate-limit-profile command

The following example creates a rate-limit profile named **hardlimit9Mb**. This rate-limit profile, when included as part of a rule in a policy list, sets a hard limit on the specified committed rate with no peak rate or peak burst ability:

```
host1(config)#ip rate-limit-profile hardlimit9Mb two-rate
host1(config-rate-limit-profile)#committed-rate 9000000
host1(config-rate-limit-profile)#committed-burst 20000
host1(config-rate-limit-profile)#committed-action transmit
host1(config-rate-limit-profile)#conformed-action drop
host1(config-rate-limit-profile)#exceeded-action drop
host1(config-rate-limit-profile)#mask-val 255
```

The following example modifies the rate-limit profile named **hardlimit9Mb** to include an exceeded action that marks the packets that exceed the peak rate. This marking action sets the DS field in the ToS byte (the six most significant bits) to the decimal value of 7 using a mask value of 0xFC:

```
host1(config)#ip rate-limit-profile hardlimit9Mb two-rate
host1(config-rate-limit-profile)#exceeded-action mark 7
host1(config-rate-limit-profile)#mask-val 252
```

To set IP precedence in the ToS byte, use the mask value of 0xE0, for visibility into the three most significant bits.

- Related Documentation**
- [Rate Limits for Interfaces Overview on page 3](#)
 - [Monitoring Policy Management Overview](#)

Setting the Peak Burst for Two-Rate Rate-Limit Profiles

For two-rate rate-limit profiles only, you can use the **peak-burst** command to set the peak burst in bytes for a rate-limit profile; range is 1–4294967295. Use to set the peak burst in milliseconds for a rate-limit profile; range is 1–10000. The no version restore the default value, 100 ms or 8192 bytes (whichever is more).

When you specify a nonzero value for the peak rate, the peak burst size is automatically calculated for a 100-ms burst as described for the **peak-rate** command. If the calculated peak burst size is less than the default value of 8192 bytes, the default value is used.



NOTE: We recommend that you do not configure a committed or peak burst size smaller than the MTU of the interface. Doing so causes large packets to be dropped even when they are transmitted at a very low rate.

- Issue the **peak-burst** command in Rate Limit Profile Configuration mode to set the peak burst in bytes:

```
host1(config-rate-limit-profile)#peak-burst 96256
```

To set the peak burst in milliseconds:

```
host1(config-rate-limit-profile)#peak-burst millisecond 1000
```

- Related Documentation**
- [peak-burst](#)

Setting a Two-Rate Rate-Limit-Profile

You can use the **rate-limit-profile two-rate** command to create a rate-limit profile and enter Rate Limit Profile Configuration mode, from which you can configure attributes for the rate-limit profile. See [Table 3 on page 21](#).



NOTE: The JunosE Software includes the layer 2 headers in the calculations it uses to enforce the rates that you specify in rate-limit profiles

Use one of the **ip**, **ipv6**, **l2tp**, or **mpls** keywords in front of the command to specify the type of rate-limit profile you want to create or modify. If you do not include one of the keywords, the router creates an IP rate-limit profile by default.

For hierarchical rate limits, do not specify the interface type, but add the **hierarchical** keyword at the end. In Parent Group Configuration Mode, associates a rate limit for a parent group. The **color-aware** keyword is only supported on hierarchical rate limits.

If you do not include a **one-rate** or **two-rate** keyword, the default is a two-rate rate-limit profile. If you enter a **rate-limit-profile** command and then type **exit**, the router creates a rate-limit profile with the default values listed in [Table 7 on page 47](#):

Table 7: Two-Rate Rate-Limit-Profile Defaults

Policy Attribute	Default Value
type	two-rate
committed-rate	0
committed-burst	8192
peak-rate	0
peak-burst	8192
committed-action	transmit
conformed-action	transmit
exceeded-action	drop
mask (IP and IPv6 rate-limit profiles)	255
exp-mask (MPLS rate-limit profiles)	7

During a software upgrade, certain values are set as follows:

- Committed burst size—Set to 8192 if it was less than that value before the upgrade
- Peak burst size—Set to 8192 if it was less than that value before the upgrade
- Peak rate—Set to 0 if it was nonzero but less than the committed rate before the upgrade



NOTE: We recommend that you do not configure a committed or peak burst size smaller than the MTU of the interface. Doing so causes large packets to be dropped even when they are transmitted at a very low rate.

- Issue the **ip rate-limit-profile** command in Global Configuration mode:

```
host1(config)#ip rate-limit-profile hardlimit9Mb two-rate
```



.....

NOTE: Commands that you issue in Rate Limit Profile Configuration mode do not take effect until you exit from that mode.

.....

- Related Documentation**
- [exp-mask](#)
 - [rate-limit-profile](#)

PART 3

Administration

- [Monitoring Tasks on page 51](#)

CHAPTER 5

Monitoring Tasks

- [Monitoring Rate-Limit Profiles on page 51](#)
- [Monitoring Control Plane Policer Information on page 53](#)

Monitoring Rate-Limit Profiles

Purpose Display information about rate-limit profiles.

Action To display information about rate-limit profiles:

```
host1#show rate-limit-profile
```

Rate Limit Profile Table

```
-----
IP Rate-Limit-Profile: rlp
  Profile Type:          one-rate
  Reference count:       0
  Committed rate:        0
  Committed burst:       8192
  Excess burst:          0
  Mask:                  255
  Committed rate action: transmit
  Conformed rate action: transmit
  Exceeded rate action:  drop

IP Rate-Limit-Profile: rlp
  Profile Type:          two-rate hierarchical
  Color-aware            no
  Reference count:       0
  Committed rate:        0
  Committed burst:       8192
  Peak rate:             0
  Peak burst:            8192
  Mask:                  255
  Committed rate action: transmit unconditional
  Conformed rate action: transmit conditional
  Exceeded rate action:  drop

L2TP Rate-Limit-Profile: L2tpRlp
  Profile Type:          two-rate
  Reference count:       0
  Committed rate:        0
  Committed burst:       8192
  Peak rate:             0
  Peak burst:            8192
  Committed rate action: transmit
  Conformed rate action: transmit
  Exceeded rate action:  drop
```

Meaning [Table 8 on page 52](#) lists the **show rate-limit-profile** command output fields.

Table 8: show rate-limit-profile Output Fields

Field Name	Field Description
Rate-Limit-Profile	Create a rate limit profile
Profile Name	Name of the rate-limit profile
Profile Type	One-rate, two-rate, or hierarchical profile
Reference count	Number of policy lists that reference this rate-limit profile
Color-aware	Color-aware action (yes or no) taken for profile
Committed rate	Target rate for the traffic, in bits per second
Committed burst	Amount of bandwidth allocated to accommodate bursty traffic, in bytes
Excess burst	Amount of bandwidth allocated to accommodate a packet in progress when the rate is in excess of the burst, in bytes
Peak rate	Amount of bandwidth allocated to accommodate traffic flow in excess of the committed rate, in bits per second
Peak burst	Amount of bandwidth allocated to accommodate bursty traffic in excess of the peak rate, in bytes
Mask	Value of mask applied to ToS byte in IP packet header
Committed rate action	Policy action (drop, transmit, or mark) taken when traffic flow does not exceed the committed rate
Conformed rate action	Policy action (drop, transmit, or mark) taken when traffic flow exceeds the committed rate but remains below the peak rate
Exceeded rate action	Policy action (drop, transmit, or mark) taken when traffic flow exceeds the peak rate

Related Documentation

- [show rate-limit-profile](#)

Monitoring Control Plane Policer Information

Purpose Display information about control plane policer for a specified protocol or all protocols.

Action To display information about control plane policer:

```
host1#show control-plane policer protocol
Burst
```

Protocol	Enabled	Rate (pps)	Size (pkts)	Packets Committed	Packets Exceeded
-----	-----	-----	-----	-----	-----
PppEchoRequest	false	50	50	0	0
PppEchoReply	false	50	50	0	0
PppEchoReplyFast	false	50	50	0	0
PppControl	false	50	50	0	0
AtmControl	false	50	50	0	0
AtmOam	false	50	50	0	0
AtmDynamicIf	false	50	50	0	0
AtmInverseArp	false	50	50	0	0
FrameRelayControl	false	50	50	0	0
FrameRelayArp	false	50	50	0	0
PppoeControl	false	50	50	0	0
PppoePppConfig	false	50	50	0	0
EthernetArp	false	50	50	0	0
EthernetArpMiss	false	50	50	0	0
EthernetLacp	false	50	50	0	0
EthernetDynamicIf	false	50	50	0	0

Meaning [Table 9 on page 53](#) lists the **show control-plane policer** command output fields.

Table 9: show control-plane policer Output Fields

Field Name	Field Description
Protocol	Name of the protocol
Enabled	True or False
Rate (pps)	Rate, in packets per second in the range 0–10000
Burst Size (pkts)	Burst size, in packets, in the range 0–10000
Packets Committed	Number of packets committed
Packets Exceeded	Number of packets exceeded

Related Documentation

- [show control-plane policer](#)

PART 4

Index

- [Index on page 57](#)

Index

A

atm commands	
atm-cell-mode.....	45

B

bandwidth management.....	22
---------------------------	----

C

color-aware configuration.....	14
committed-action command.....	34
committed-burst command.....	35
conventions	
notice icons.....	ix
text and syntax.....	x
customer support.....	xi
contacting JTAC.....	xi

D

documentation set	
comments on.....	xi

E

exceeded-action command.....	37
excess-burst command.....	37
exp-mask command.....	37

M

manuals	
comments on.....	xi
mark rule	
with rate-limit rule.....	34

N

notice icons.....	ix
-------------------	----

O

one-rate rate-limit profile.....	23
----------------------------------	----

P

peak-burst command.....	46
peak-rate command.....	37

percent-based rates.....	15
--------------------------	----

policy management

committed burst calculation.....	35, 37
congestion management.....	22
creating a one-rate rate-limit profile.....	39
creating a two-rate rate-limit profile.....	45
modifying a one-rate rate-limit profile.....	39
modifying a two-rate rate-limit profile.....	45
one-rate rate-limit profile.....	23
policy actions and rate-limit profiles.....	20
rate limiting traffic flows.....	24
rate-limit profile actions.....	20
rate-limit profile attributes.....	20
rate-limit profile calculations.....	42
rate-limit profile defaults.....	42, 46

policy parameter

considerations.....	17
quick configuration.....	34
reference-rate.....	15

R

rate limiting

aggregate traffic flows.....	24
for interfaces.....	3
hierarchical.....	4
individual traffic flows.....	24
rate-limit hierarchies.....	4
classifier groups.....	4
color-aware configuration.....	14
rate-limit actions.....	4
rate-limit profiles.....	4
rate-limit rules.....	4

rate-limit profiles

attributes.....	20
burst size.....	15
calculations	42
configuration procedure.....	29
creating.....	39, 45
default values.....	42, 46
modifying.....	39, 45
percent-based rates.....	15
policy actions.....	20
rates.....	15

rate-limit rule

with mark rule.....	34
---------------------	----

rate-limit-profile one-rate command.....	42
--	----

rate-limit-profile two-rate command.....	46
--	----

rate-limiting SRP traffic flows.....	38
--------------------------------------	----

S

show commands

show control-plane policer	53
show parent-group.....	51
show rate-limit-profile	51
single-rate rate limit profile.....	19
support, technical	See technical support

T

TCP-friendly one-rate rate-limit profile.....	40
technical support	
contacting JTAC.....	xi
text and syntax conventions.....	x
two-rate rate-limit profile.....	24