

Network Working Group
Request for Comments: 2697
Category: Informational

J. Heinanen
Telia Finland
R. Guerin
University of Pennsylvania
September 1999

A Single Rate Three Color Marker

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.


Abstract

This document defines a Single Rate Three Color Marker (srTCM), which can be used as component in a Diffserv traffic conditioner [RFC2475, RFC2474]. The srTCM meters a traffic stream and marks its packets according to three traffic parameters, Committed Information Rate (CIR), Committed Burst Size (CBS), and Excess Burst Size (EBS), to be either green, yellow, or red. A packet is marked green if it doesn't exceed the CBS, yellow if it does exceed the CBS, but not the EBS, and red otherwise.

1. Introduction

The Single Rate Three Color Marker (srTCM) meters an IP packet stream and marks its packets either green, yellow, or red. Marking is based on a Committed Information Rate (CIR) and two associated burst sizes, a Committed Burst Size (CBS) and an Excess Burst Size (EBS). A packet is marked green if it doesn't exceed the CBS, yellow if it does exceed the CBS, but not the EBS, and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility.

The Meter meters each packet and passes the packet and the metering result to the Marker:



The Meter operates in one of two modes. In the Color-Blind mode, the Meter assumes that the packet stream is uncolored. In the Color-Aware mode the Meter assumes that some preceding entity has pre-colored the incoming packet stream so that each packet is either green, yellow, or red. The details of the pre-coloring process, including handling of error scenarios, and how the Meter determines the color of a pre-colored packet are DS domain specific and outside the scope of this document.

The Marker (re)colors an IP packet according to the results of the Meter. The color is coded in the DS field [RFC2474] of the packet in a PHB specific manner (see section 4 for an example).

A companion document [RFC2698] describes another three color marker, called a Two Rate Three Color Marker (trTCM), where packets are marked based on two rates and two burst sizes.

2. Configuration

The srTCM is configured by setting its mode and by assigning values to three traffic parameters: a Committed Information Rate (CIR), a Committed Burst Size (CBS), and an Excess Burst Size (EBS).

The CIR is measured in bytes of IP packets per second, i.e., it includes the IP header, but not link specific headers.

The CBS and the EBS are measured in bytes. The CBS and EBS must be configured so that at least one of them is larger than 0. It is recommended that when the value of the CBS or the EBS is larger than 0, it is larger than or equal to the size of the largest possible IP packet in the stream.

3. Metering

The behavior of the Meter is specified in terms of its mode and two token buckets, C and E, which both share the common rate CIR. The maximum size of the token bucket C is CBS and the maximum size of the token bucket E is EBS.

The token buckets C and E are initially (at time 0) full, i.e., the token count $Tc(0) = CBS$ and the token count $Te(0) = EBS$. Thereafter, the token counts Tc and Te are updated CIR times per second as follows:

- o If Tc is less than CBS, Tc is incremented by one, else
- o if Te is less than EBS, Te is incremented by one, else
- o neither Tc nor Te is incremented.

When a packet of size B bytes arrives at time t, the following happens if the srTCM is configured to operate in the Color-Blind mode:

- o If $Tc(t) - B \geq 0$, the packet is green and Tc is decremented by B down to the minimum value of 0, else
- o if $Te(t) - B \geq 0$, the packet is yellow and Te is decremented by B down to the minimum value of 0, else
- o the packet is red and neither Tc nor Te is decremented.

When a packet of size B bytes arrives at time t, the following happens if the srTCM is configured to operate in the Color-Aware mode:

- o If the packet has been precolored as green and $Tc(t) - B \geq 0$, the packet is green and Tc is decremented by B down to the minimum value of 0, else
- o If the packet has been precolored as green or yellow and if $Te(t) - B \geq 0$, the packet is yellow and Te is decremented by B down to the minimum value of 0, else
- o the packet is red and neither Tc nor Te is decremented.

Note that according to the above rules, marking of a packet with a given color requires that there be enough tokens of that color to accommodate the entire packet. Other marking policies are clearly possible. The above policy was chosen in order to guarantee a

deterministic behavior where the volume of green packets is never smaller than what has been determined by the CIR and CBS, i.e., tokens of a given color are always spent on packets of that color.

The actual implementation of a Meter doesn't need to be modeled according to the above formal specification.

4. Marking

The Marker reflects the metering result by setting the DS field of the packet to a particular codepoint. In case of the AF PHB [RFC2597], the color can be coded as the drop precedence of the packet.

5. Service Example

The srTCM can be used to mark a packet stream in a service, where different, decreasing levels of assurances (either absolute or relative) are given to packets which are green, yellow, or red. For example, a service may discard all red packets, because they exceeded both the committed and excess burst sizes, forward yellow packets as best effort, and forward green packets with a low drop probability.

6. Security Considerations

The srTCM has no known security concerns.

7. References

[RFC2698] Heinanen, J. and R. Guerin, "A Two Rate Three Color Marker", RFC 2698, September 1999.

[RFC2597] Heinanen, J., Baker, F., Weiss, W. and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.

[RFC2474] Nichols, K., Blake, S., Baker, F. and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

8. Authors' Addresses

Juha Heinanen
Telia Finland, Inc.
Myyrmaentie 2
01600 Vantaa, Finland

EMail: jh@telia.fi

Roch Guerin
University of Pennsylvania
Department of Electrical Engineering, Rm 376 GRW
200 South 33rd Street
Philadelphia, PA 19104

EMail: guerin@ee.upenn.edu

9. Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

