Network Working Group
Request for Comments: 4115
Category: Informational

O. Aboul-Magd
S. Rabie
Nortel Networks

July 2005

A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of in-Profile Traffic

Status of This Memo

Copyright Notice

IESG Note

Abstract

This document describes a two-rate, three-color marker that has been in use for data services including Frame Relay services. This marker can be used for metering per-flow traffic in the emerging IP and L2 VPN services. The marker defined here is different from previously defined markers in the handling of the in-profile traffic. Furthermore, this marker doesn't impose peak-rate shaping requirements on customer edge (CE) devices.

1. Introduction

The differentiated service defines a quality-of-service (QoS) architecture for the Internet [RFC2475]. Two integral components of this architecture are traffic metering and marking. This document describes a two-rate, three-color metering/marker algorithm that is

suitable for the differentiated service applications such as IP and L2 VPNs. This algorithm has been in use for data services including Frame Relay Service.

The metering/marker defined here is different from those in [RFC2697] and [RFC2698]. It is different from [RFC2697] in that it is a two-rate, three-color marker. In contrast, [RFC2697] is a single-rate marker. It is different from [RFC2698] in the way its parameters are defined, which allows a better handling of in-profile traffic for predominant service scenarios over a wider range of traffic parameters.

Furthermore, the algorithm described here eliminates the need for the CE to shape its traffic to a certain peak information rate (PIR), as might be the case for the marker defined in [RFC2698] when the value for the peak burst size (PBS) is smaller than that for the committed burst size (CBS).

The marker described here operates for both color-blind and color-aware modes, as defined in [RFC2698].

2. Configuration

The operation of the marker is described by two rate values. The committed information rate (CIR) and the excess information rate (EIR). CIR and EIR define the token generation rate of a token bucket with size that is equal to committed burst size (CBS) and excess burst size (EBS), respectively.

The CBS and EBS are measured in bytes and must configure to be greater than the expected maximum length of the incoming PDU. The CIR and EIR are both measured in bits/s. The CIR and EIR can be set independently of each other. Alternatively, the CIR and EIR can be linked together by defining a burst duration parameter, T, where T=CBS/CIR=EBS/EIR.

3. Metering and Marking

The behavior of the meter is defined in terms of its mode and two token buckets, C and E, with the rates, CIR and EIR, respectively, and maximum sizes CBS and EBS.

The token buckets C and E are initially (at time 0) full; i.e., the token count Tc(0) = CBS and Te(0) = EBS. Thereafter, the token count Tc is incremented by one CIR times per second (up to CBS), and the token count Te is incremented by one EIR times per second (up to EBS).

In the color-aware operation, it is assumed that the algorithm can recognize the color of the incoming packet (green, yellow, or red). The color-aware operation of the metering is described below.

When a green packet of size B arrives at time t, then

- if Tc(t)- B > 0, the packet is green, and Tc(t) is decremented by B; else
- if Te(t)- B > 0, the packet is yellow, and Te(t) is decremented by B; else
- the packet is red.

When a yellow packet of size B arrives at time t, then

- if Te(t)- B > 0, the packet is yellow, and Te(t) is decremented by B; else
- the packet is red.

Incoming red packets are not tested against any of the two token buckets and remain red.

In the color-blind operation, the meter assumes that all incoming packets are green. The operation of the meter is similar to that in the color-aware operation for green packets.

The salient feature of the algorithm described above is that traffic within the defined CIR is colored green directly, without the need to pass additional conformance tests. This feature is the main differentiator of this algorithm from that described in [RFC2698], where traffic is marked green after it passes two conformance tests (those for PIR and CIR). In either color-blind or color-aware mode, the need to pass two conformance tests could result in packets being dropped at the PIR token bucket even though they are perfectly within their CIR (in-profile traffic). Furthermore, in the color-aware mode of operation, the need to pass two conformance tests could make yellow traffic starve incoming in-profile green packets.

The operation of the algorithm is illustrated in the flow chart below:

Figure 1: Traffic Metering/Marking Algorithm

In Figure 1, we have X(t-) and X(t+) to indicate the value of a parameter X right before and right after time t.

4. Service Scenario

The described marker can be used to mark an IP packet stream in a service where different, decreasing levels of assurances (either absolute or relative) are given to packets that are green, yellow, or red. For example, a service may discard all red packets because they exceeded the service rates, forward yellow packets as best effort, and forward green packets with low drop probability. The marker could also be used for metering L2 VPN services such as the emerging Ethernet transport over IP networks.

5. Security Considerations

Security issues resulting from this document are similar to those mentioned in [RFC2697] and [RFC2698].

6. Informative References

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

[RFC2697] Heinanen, J. and R. Guerin, "A Single Rate Three Color Marker", RFC 2697, September 1999.

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

[RFC3932] Alvestrand, H., "The IESG and RFC Editor Documents: Procedures", BCP 92, RFC 3932, October 2004.

Authors' Addresses

Osama Aboul-Magd
Nortel Networks
3500 Carling Ave
Ottawa, ON K2H8E9
EMail: osama@nortel.com

Sameh Rabie
Nortel Networks
3500 Carling Ave
Ottawa, ON K2H8E9
EMail: rabie@nortel.com

Full Copyright Statement

Intellectual Property

Acknowledgement

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