Internet Engineering Task Force
INTERNET DRAFT
Expires February 2000

Longsong Lin
Jeffrey Lo
NEC USA Inc.

Fang-Ching Ou
National Yunlin University

August,1999

A Generic Traffic Conditioner
<draft-lin-diffserv-gtc-01.txt>

Status of this Memo

Abstract

This document defines a Generic Traffic Conditioner (GTC) for Diffserv-capable nodes [RFC2475, RFC2474]. The GTC employs a token bucket to meter the traffic stream, and a shaping buffer to store out-of-profile packets of a behavior aggregate in order to enforce conformance to the traffic profile. The GTC employs a policer to decide whether to shape, drop, or re-mark the out-of-profile packets. The GTC marks each outgoing packet with a Diffserv codepoint (DSCP) according to the traffic conditioning result in accordance with the Traffic Conditioning Specification. By enabling the components and configuring them with associated traffic parameters, the GTC can be used to condition the traffics for EF PHB [Jacobson], AF PHBs [Heinanen], as well as Class Selector PHBs [RFC 2474].

1. Introduction

The Generic Traffic Conditioner (GTC) consists of five functional elements: meter, policer, dropper, marker, and shaper as shown in Figure 1. The meter measures an IP packet stream and recognizes the packets as either in-profile or out-of-profile according to the traffic profile in the Traffic Conditioning Specification (TCS). The result of metering is reflected in the marker where the in-profile packets are marked with a corresponding DSCP. Usually, marking takes place at DS boundary node that connects one DS domain to another. However, the in-profile packets could bypass the marker to a fowarding PHB if marking is not necessary. The meter then hands over the out-of-profile packets to the policer where the decision is made as to shape, drop or re-mark the packets in accordance with the provisioning policy in the TCS. The shaper could delay the out-of-profile packets which will be fedback as an input to the meter for re-metering against the same traffic profile. The dropper discards out-of-profile packets. The marker could re-mark the out-of-profile packets with a different corresponding DSCP according to the result of policing.

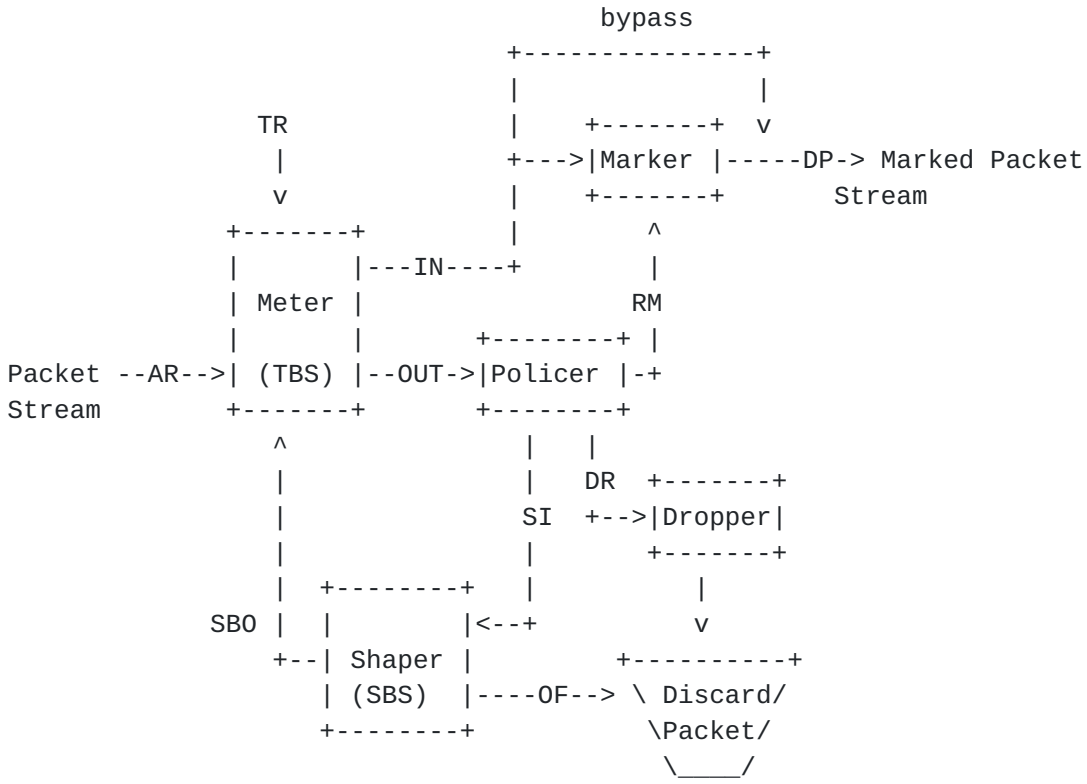

Figure 1. Generic Traffic Conditioner (GTC)

2. Configuration and Monitoring

The GTC is configured by enabling the selected components and by assigning values to three traffic parameters: Token Rate(TR), Token Bucket Size (TBS), and an Shaping Buffer Size (SBS).

The GTC MAY be configured to serve traffic contracted at different PHBs. For example, it may function as a TC for EF PHB by enabling a meter, a policer, a shaper, a dropper, and a marker, i.e., with all components enabled. Next, it may function as a TC for AF PHB by enabling a meter, a policer, dropper, and a marker.

It MUST be possible to monitor over certain time interval the current Token Bucket Occupancy (TBO), Shaping Buffer Occupancy (SBO), the number of bytes that have arrived at the shaper (SI), the in-profile packets in bytes (IN), the out-of-profile packets in bytes (OUT), the number of bytes of out-of-profile packets that are dropped (DR), the number of bytes that have been re-marked (RM) and the number of bytes that overflow the shaping buffer (OF).

The token rate (TR) is measured in bytes of IP packets per unit time. It is assumed for simplicity that the unit time be one second and that all corresponding operations of the components in GTC MUST be completed within unit time. The TBS, TBO, SBS and SBO are measured in bytes. Each packet includes the IP header, but not link specific headers.

3. Minimum Functionality

The TBS MUST be configured to be larger than 0 and the SBS MAY be configured larger than or equal to 0. It is RECOMMENDED that the value of the TBS or the SBS, if larger than 0, be larger than or equal to the size of the largest possible IP packet in the stream. The SBS, if larger than 0, and TBS MUST be configured at least larger than maximum transimssion unit (MTU).

The TBS SHOULD be configured with respect to the requirements of the service and the jitter and delay due to the configuration should not cause any violation of these requirements. The SBS SHOULD be configured with respect to the TBS and TR. It is possible to configure the SBS as up to two bandwidth-delay products. However, since shaping often introduces jitter and delay for packets being shaped, it is RECOMMENDED that the SBS be configured according to the requirements of the service and the jitter and delay incurred due to shaping should not result in any violation of these requirements.

The output of the GTC SHOULD average TR and MUST not exceed TBS+TR*t when measured over certain time interval t equal to or longer than

the time it takes to send an output link MTU sized packet at TR.

4. Algorithm

The behavior of the GTC is specified in terms of two decision elements (metering and policing) and three respective action elements (marking, shaping, and dropping).

The meter measures the incoming packets against the measurement specified in the TCS, e.g., the inter-arrival rate, to determine whether the packets conform to the profiled measurement (in-profile) or otherwise (out-of-profile). The meter SHOULD not change the behavior of a packet stream but only identify the conformance characteristics of each packet. While the in-profile packets will be marked with a corresponding DSCP, the out-of-profile packets are subjected to policing.

The policer controls the subsequent treatment of each out-of-profile packet, subject it to shaping, re-marking, or dropping. The decision as to which action to take should consider the delay, jitter or loss requirement of the packet stream.

An out-of-profile packet may be enqueued to a shaper which can regulate the packet to conform to the traffic profile via re-measuring. The size of shaping buffer is constrained by the delay requirement of the packet stream and shaping is possible only if the buffer is not overflowed. An out-of-profile packet may be transformed to other classes of service by re-marking it with a different DSCP; otherwise, it may be silently discarded by a dropper. The decision as to re-marking or dropping depends on the provisioning policy at the node.

The marker marks (re-marks) packets according to the results of metering (policing). A diagrammatic representation of the algorithm is given in Figure 2.

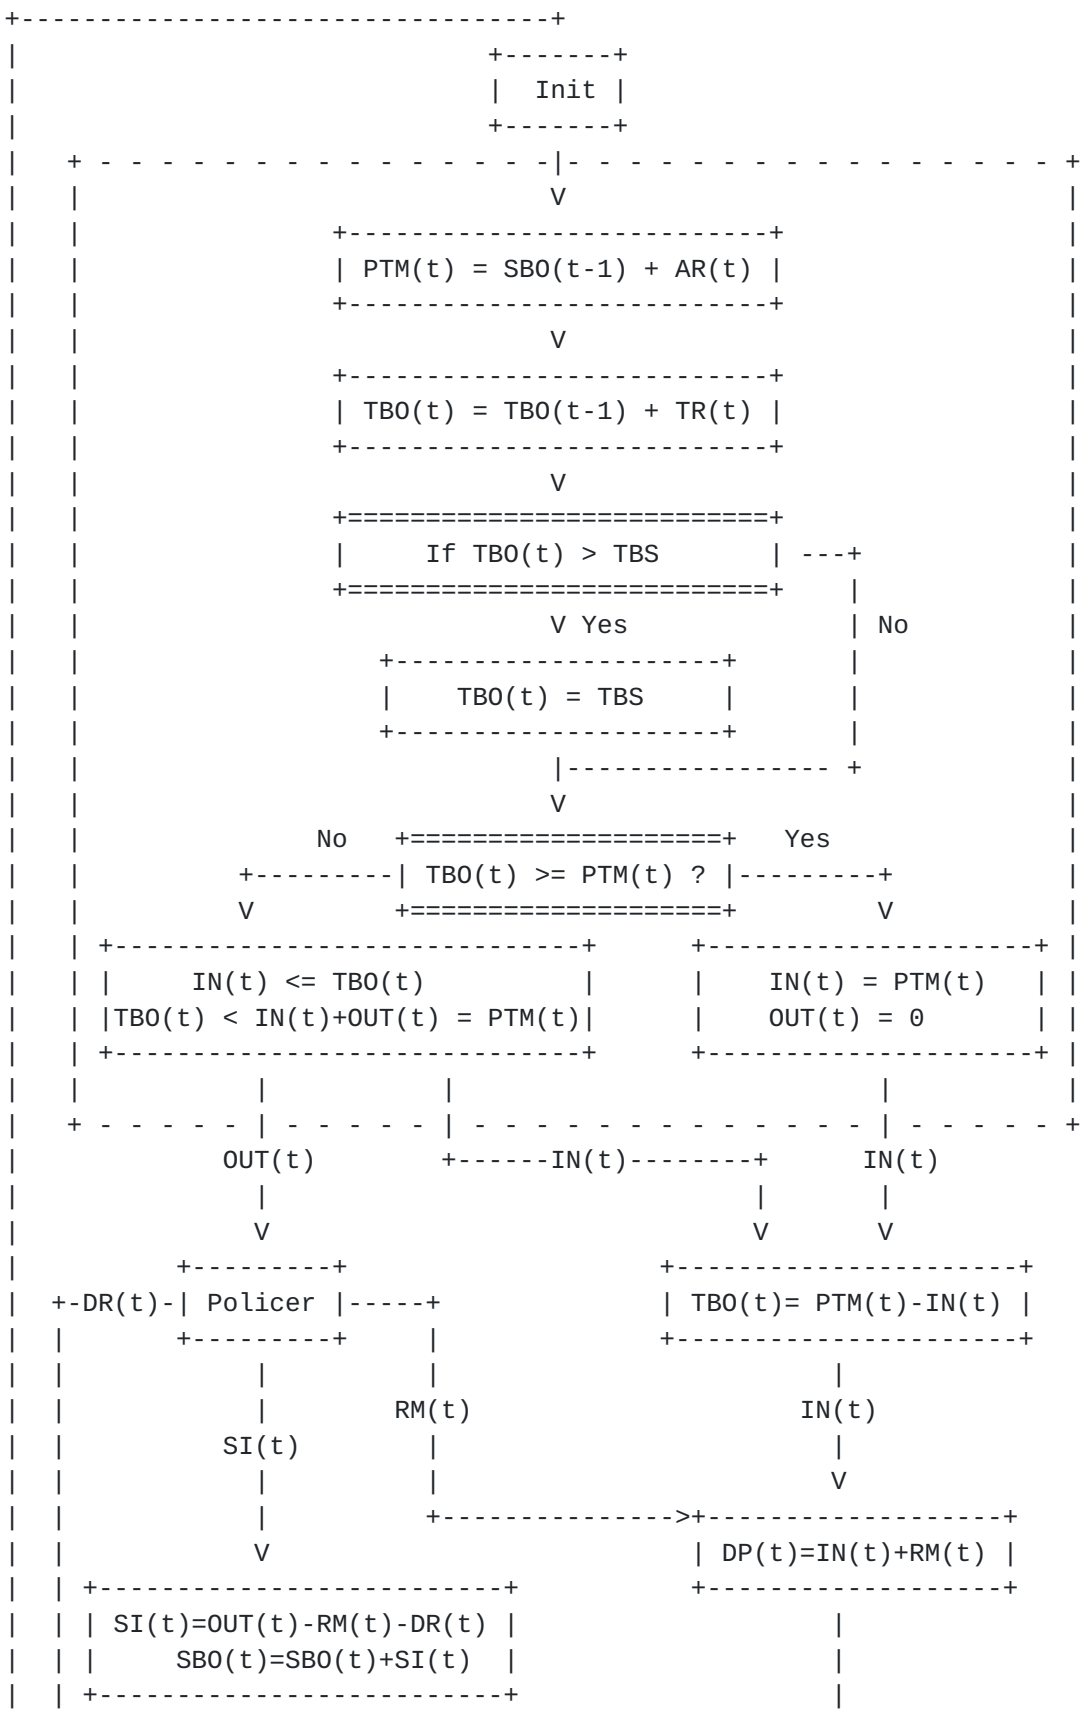
Init
PTM(t) = SBO(t-1) + AR(t)
TBO(t) = TBO(t-1) + TR(t)
If TBO(t) > TBS
Yes
No
TBO(t) = TBS
No
TBO(t) >= PTM(t) ?
Yes
IN(t) <= TBO(t)
TBO(t) < IN(t)+OUT(t) = PTM(t)
IN(t) = PTM(t)
OUT(t) = 0
OUT(t)
IN(t)
IN(t)
DR(t)
Policer
RM(t)
IN(t)
SI(t)
TBO(t)= PTM(t)-IN(t)
DP(t)=IN(t)+RM(t)
SI(t)=OUT(t)-RM(t)-DR(t)
SBO(t)=SBO(t)+SI(t)

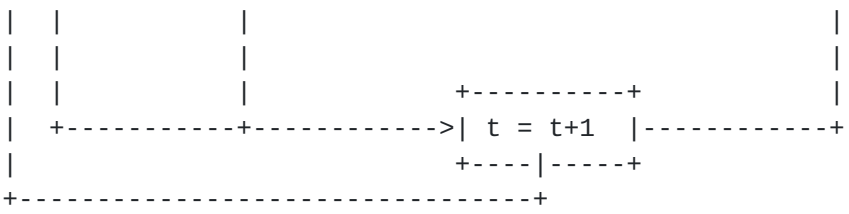

Figure 2. GTC Algorithm

The token bucket is initially (at time 0) full, i.e., the token bucket occupancy TBO(0) = TBS, while the shaping buffer is initially empty, i.e., the shaping buffer occupancy SBO(0) = 0.

The TBO is increased by TR bytes per unit time. Hence the total available token for the data stream at time t is

- TBO(t) = TBS, for t=0.
- TBO(t) = TBO(t-1)+TR(t), for t>0.
- If TBO(t) > TBS, then TBO(t) =TBS.

The packets to be measured (PTM) by token bucket come from two sources: one from the input arrival at the traffic conditioner (AR) and the other from the feedback from the output of shaping buffer (SBO), i.e.,

- PTM(t) = SBO(t-1)+AR(t).

To prevent packet reordering, the packets in the shaping buffer MUST be measured before any new packet arriving at the GTC.

If TBO(t) < PTM(t), then the total number of bytes in the packets that are in-profile (IN) will be IN(t) <= TBO(t), and the total number of bytes in the packets that are out-of-profile (OUT) should satisfy TBO(t) < IN(t)+OUT(t) = PTM(t).

If TBO(t) >= PTM(t), then all the packets are in-profile, i.e. IN(t)= PTM(t) and OUT(t)=0.

The OUT(t) in any unit time may be made up of multiple packets. The policer operates on each of these packets in turn. It decides how each of the out-of-profile packets is treated depending on the present occupancy of the shaping buffer as well as service policy. Packets whose addition will not results in violation of expected per-hop behavior may be submitted to the shaping buffer for re-measuring; otherwise, they should be dropped. Packets not dropped will be enqueued for shaping or remarked depending on the service policy. For each out-of- profile packet i with length Pi,

o If SBO(t)+Pi > SBS, then
drop the packet, DR(t) += Pi,
o Else
shape the packet, SI(t) += Pi, Or
remark the packet, RM(t) += Pi.

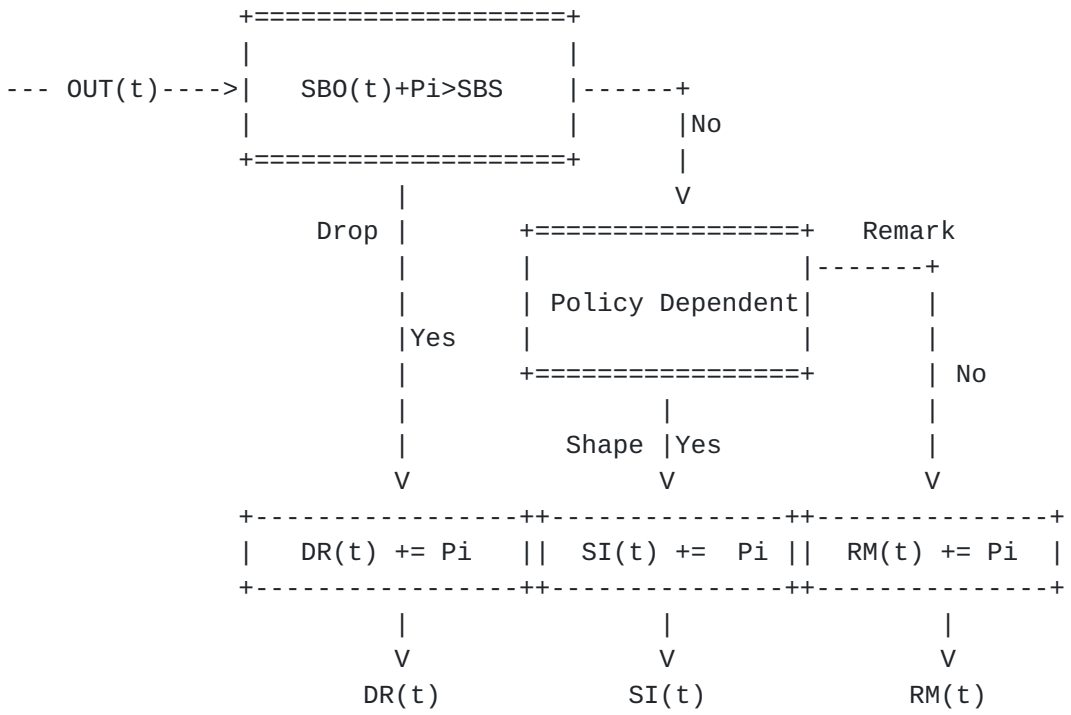

Figure 3. Policer

Thereafter, the token bucket and shaping buffer will be updated as follows.

For packets in IN(t),
- o TBO(t) = TBO(t)-IN(t).

For packets in OUT(t),
- o SBO(t) = SBO(t-1) + SI(t), where SI(t)=OUT(t)-RM(t)-DR(t).
- o If SBO(t) > SBS, then
  discard the overflow packets (OF) and
  SBO(t) = SBO(t)- OF(t).

5. Re-marking

The marker marks the in-profile packets according to traffic profile defined in the Traffic Conditioning Specification (TCS). The marker will distinguish whether or not the current boundary node is a first-hop router [RFC2475] or an interior node. If it is a first-hop router , it may mark the unmarked packets or re-mark previously marked packets to an effective codepoint.

The marker also re-marks packets that are out-of-profile subject to the service policy. Re-marking may result in either change of PHB group or change of drop preference in a service class. For example, re-marking could demote a packet from an AF PHB to the Default best-effort PHB. As another example, re-marking may change the drop precedence of the packet in AF PHB class [Heinanen]. Re-marking MUST not cause packet reordering.

6. Example Services

The GTC can be used to serve EF PHB and AF PHB. The GTC conditions packet aggregate via policing and shaping so that its arrival rate to a PHB at any node is always less than that node's configured minimum departure rate for the PHB. This property ensures the traffic conditioning requirement for EF PHB. EF flow packets see sufficient tokens present will have their EF codepoints marked. When tokens are not sufficient, EF packets will be held in the shaping buffer until tokens arrive. If an EF flow bursts enough to overflow the shapping buffer, its packet will be dropped.

For an AF PHB class, in-profile packets and out-of-profile packets are distinguished by marking them with different AF codepoints. For examples, the in-profile packet can be marked as AFx1 and the out-of-profile packets can be policed to be marked as AFx(2,3), omitting shaping and dropping. The GTC may also shape the out-of-profile AF packets if necessary.

The GTC can support CS DSCP with preferential forwarding marking. The relative order of the CS DSCPs can be implemented by using the policer and remarker to transform from one CS DSCP to another. The

probability of timely forwarding can be enforced by GTC token bucket mechanism as well as by the CS-compliant PHB mechanisms. As an extreme case, packet dropped by a GTC policer is considered to be untimely forwarding. A GTC meter should be configured with a minimum rate for the CS DSCP behavior agregates in accordance with the CS-compliant PBH and the default PHB.

7. Assumptions

The GTC has no known security concerns and no a priori information for input traffic.

8. Discussion

The arrangement of the components of the GTC in this document is by no means the only solution to provide conditioning for diffserv traffic. A simple configuration for providing a particular service is possible. For example, a cascaded configuration of a meter with a marker is sufficient for supporting AF PHB.

The components in the GTC is a generic representation of funcationality; each component can be extended to cover more complex functions. For example, the meter could output multiple levels of conformance information, instead of two levels of IN or OUT of profile information. A multi-color marker is an example of the extension.

It is noted that any other components which offer new functionality for out-of-profile packets can easily be embedded in the GTC as long as it is configured appropriately. For instance, a spacer can be employed on non-conforming traffic to reduce the delay variation. Finally, it is noted that the GTC can operate in accordance to the periodical sampling times or to the events by packet arrivals.

9. References

[Heinanen] J. Heinanen, et al., Assured Forwarding PHB Group. Internet draft draft-ietf-diffserv-af-06.txt, February 1999.

[Nichols] K. Nichols and B. Carpenter, Format for Diffserv Working Group Traffic Conditioner Drafts. Internet draft draft-ietf-diffserv-traffcon-format-00.txt, February 1999.

[RFC2474] K. Nichols, et al., Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. RFC 2474, December 1998.

[RFC2475] S. Blake, et al., An Architecture for Differentiated

Services. RFC 2475, December 1998.

[Heinanen1] J. Heinanen, et al. A Three Color Marker. Internet draft draft-heinanen-diffserv-tcm-01.txt, February 1999.

[Jacobson] V.Jacobson, et al., Expedited Forwarding Per Hop Behavior Internet Draft draft-ietf-diffserv-phb-ef-02.txt, February 1999.

10. Author Address

Longsong Lin
NEC USA, Inc.
110 RIO ROBLES
SAN JOSE, CA 95134, USA
Email : lin@ccrl.sj.nec.com
Tel : (408) 943 3032

Fang-Ching Ou
Dept. of Electronic and Information Engineering
National Yunlin University of Science and Technology
123 University Rd. Section 3, Touliu City, Yunlin, Taiwan
Email : ofcaltos@ms14.hinet.net

Jeffrey Lo
NEC USA, Inc.
110 RIO ROBLES
SAN JOSE, CA 95134, USA
Email : jlo@ccrl.sj.nec.com
Tel : (408) 943 3033