Internet Engineering Task Force
Victor Firoiu,
Bay Networks/Nortel Networks
Alessio Casati,
Lucent Technologies

INTERNET DRAFT
Document: <draft-firoiu-diffserv-af-amend-00.txt>
November, 1998
Expires: May 1999

Amendments to the Assured Forwarding PHB Group

Status of this Memo

1. Abstract

This note was motivated by the ongoing discussion in the Differentiated Services Working Group regarding the definition of the Assured Forwarding (AF) Per-Hop Behavior (PHB) Group. We consider two issues with the current proposal for AF PHB Group in [Heinanen]: the recommendation for AF discard mechanism, and the definition of AF classes. We discuss the implications of the definitions and recommendations that have raised comments and propose alternative texts.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119].

3. Introduction

We are concerned with two aspects in the current AF PHB proposal in [Heinanen]. The first issue regards the recommendation for AF discard mechanism. The second issue relates to the definition of AF classes. In Section 4, we discuss the first issue, and in Section 5, the second.

4. Recommendation for AF discard mechanism

[Heinanen] Section 4 has the following text:

Text 1. "Inconsistent discard behaviors lead to inconsistent end-to-end service semantics. It is RECOMMENDED that the discard mechanism is based on a RED-like [Floyd] algorithm."

Consistency of local packet treatment is clearly needed for providing a defined end-to-end service. However, a uniform implementation is neither necessary nor sufficient for providing such consistent per-hop behavior. Specifically, it is possible to provide consistent discard behaviors with implementations different than RED-based. Moreover, recommending RED-based implementations does not guarantee consistent discard behavior, since it is possible to construct pathological configurations of RED.

We consider that, in order to ensure consistency, we need to recommend the discard behavior, and not an implementation. We propose to define the AF drop preference behavior as follows, replacing Text 1:

"Inconsistent discard behaviors lead to inconsistent end-to-end service semantics. It is RECOMMENDED that the discard mechanism be based on active queue management algorithms. An example of such a discard mechanism is a RED-like [Floyd] algorithm with three configurable levels of drop precedence."

5. Definition of AF classes

The second issue relates to the definition of AF classes. The following text in [Heinanen] is the most relevant to this issue.

Text 2. "A DS node MUST allocate forwarding resources (buffer space and bandwidth) to AF classes so that, under reasonable operating conditions and traffic loads, packets of an AF class x do not have higher probability of timely forwarding than packets of an AF class y if x < y."

Timely forwarding is thus the criterion differentiating the four AF classes. Unfortunately, this is not a well-defined criterion, and can give way to different and inconsistent interpretations. For example, what is a conformance test for "packets of AF class 1 do not have higher probability of timely forwarding than packets of AF class 2"? Is timely delivery the average (or minimum, or maximum) packet forwarding time at a DS node? An end-to-end service constructed from DS nodes using different interpretations can be inconsistent or unpredictable.

We consider that this issue is essential for the definition of AF PHB Group, and that more discussion is needed in the Working Group to clarify it.

As a basis for discussion, we propose to define the timeliness of packet forwarding in an AF Class at a DS node to be the average forwarding time of non-dropped packets. The forwarding time is the time between the input and output of a packet at the DS node. The averaging time interval is to be defined or may be a differentiation parameter among different providers.

Observe that a precise definition of timely forwarding does not make the AF a quantitative PHB. For example "average forwarding time of Class 1 smaller than average forwarding time of Class 2" is a qualitative characterization.

Another issue is that AF class differentiation through forwarding time can be realized in many different ways. For example, one way to provide small average forwarding time for an AF class is a combination of low average load and large amount of resources. Another way is to have an aggressive drop policy on low priority packets, thus a small average queue length and ultimately small average queuing delay. Therefore, we believe that it is better to differentiate AF classes by a behavioral attribute, the average forwarding time, leaving configuration details open to innovation.

In conclusion, we propose to define the forwarding behavior provided to packets belonging to different AF classes in the following way, replacing Text 2:

"In a DS node, under normal operating conditions and traffic loads, packets of AF class x do not have smaller forwarding time (delay) than packets of AF class y if x < y, the averaging being over the same time interval."

6. References

[Heinanen], Heinanen, J, Baker, F, Weiss, W, Wroclawski, J, "Assured Forwarding PHB Group", draft-ietf-diffserv-af-03.txt, November, 1998.

[Floyd], Floyd, S., and Jacobson, V., Random Early Detection gateways for Congestion Avoidance. IEEE/ACM Transactions on Networking, Volume 1, Number 4, August 1993, pp. 397-413.

[RFC-2119], Bradner, S, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, Harvard University, March 1997

7. Author's Addresses

Victor Firoiu
Bay Networks / Nortel Networks
3 Federal St
Billerica MA, 01821
Phone: +1 (978) 916-4354
Email: vfiroiu@baynetworks.com

Alessio Casati
Lucent Technologies
Sigma
Windmill Hill Business Park
Swindon
Wiltshire SN5 6PP
United Kingdom
Phone: +44 (1793) 88-3861
Email: acasati@lucent.com