INTERNET-DRAFT
draft-brunner-diffserv-pdb-one2any-ar-00.txt
Expires August 2001

M. Brunner
A. Banchs
S. Tartarelli
H. Pan
NEC Corporation

February 2001

An one-to-any Assured Rate Per-Domain Behavior for Differentiated Services

<draft-brunner-diffserv-pdb-one2any-ar-00.txt>

Status of this Memo

Abstract

This document defines a Per-Domain Behavior (PDB) called one-to-any Assured Rate. The PDB is useful to implement services in a Differentiate Services domain, which need an assured rate. The assurance is given with a certain well-defined probability for traffic using this PDB. However, no delay and no jitter guarantees are provided.

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.

Table of Contents

1. Introduction

This document defines a differentiated services Per-Domain Behavior (PDB) suitable for traffic that requires rate assurance but does not require delay and jitter bounds.

The document defines a Per-Domain Behavior similar to the one given in [AR-PDB]. However, we address the one-to-any case only. The PDB provides an assured rate to users of the network with a well-defined probability for traffic conforming to a negotiated rate. No quality assurance can be given for traffic above the assured rate.

The assured rate PDB defined in [AR-PDB] may work well for one-to-one scenarios, since an admission rule for a one-to-one PDB case is relatively simple. In the one-to-one case, the destination of the traffic aggregate is known. Assuming static routing, the links involved in transmitting the aggregate can be identified and capacity in those links can be reserved. Note that sending at full rate assured rate traffic might use all the service class' capacity.

But providing an assured rate with almost no drops can be too expensive for the one-to-any case. Additionally, an admission rule for the one-to-any case is more difficult to derive. With a one-to-any PDB, the links involved with the traffic aggregate vary depending on the destination to which the user is sending at some point in time. Reserving capacity in all possible links may be too expensive, since this would be a hard limit to the amount of assured rate traffic that could be accepted. Note that in this case, assured rate traffic could at most utilize only a small portion of the total

service class' capacity, resulting in a low efficiency. However, if capacity is not reserved for all possible links, there will be a certain probability that the rate assurances are not met. We argue that having a certain probability of not meeting the rates assurances is necessary in order to be able to accept a reasonable

amount of assured rate traffic in the network for the one-to-any PDB. Therefore we propose a new PDB which considers the possibility that assured rates commitments are not met with a certain probability.

The one-to-any AR PDB may be used to build services where the destination is not known in advance, but a certain sending traffic to any location is still required. So the basic service, which may be built over this PDB, is one for sending data with an assured rate to any location. A key feature of this kind of services is that the destination of a connection can be any egress point of the ISP's DS domain.

Note that the edge conditioning rules and most of the parameters and attributes of the PDB are the same as specified in [AR-PDB]. Additional to the already proposed draft, we add an admission rule for the one-to-any case.

2. Description of the one-to-any assured rate PDB

We define a one-to-any traffic aggregate as traffic arriving at one DS domain ingress router and leaving the DS domain at any egress router. With assured rate, we refer to a rate, which is assured with a certain probability. However, no delay and jitter guarantees for a traffic aggregate are given.

This PDB assures that traffic conforming to an assured rate will not be dropped with a specified probability. A user may send traffic that excesses the assured rate, but no guarantees are given to obtain unused bandwidth in the network.

This PDB is referred to as the one-to-any Assured Rate (AR) PDB and is defined in accordance with the guidelines in [PDBDEF].

3. Applicability Statement

This document does not restrict the PDB to any particular application or traffic type. Regardless of the traffic model, the traffic aggregate will get the assured rate.

However, the PDB only applies in large enough networks with many users, where statistical multiplexing in terms of user behavior works. We assume statistical description of user aggregated behavior in terms of what egress router the traffic of the aggregate flows to.

4. Technical Specification

The rule specification for this PDB consist of three parts:

1. A set of edge rules to classify packets arriving at the domain ingress into a traffic aggregate, perform metering/policing on the aggregate and associate a packet marking with the aggregate.

2. Per-node PHB treatment for the traffic from the ingress to the egress.

3. Admission rules that specify whether a new PDB instance (the PDB attributes filled with values) is accepted. Note that we rely on admission control to assure the rate for provisioned traffic aggregates.

4.1. Edge Rules

As packets enter the domain they will be classified into a traffic aggregate based on the specified classifier rules at the domain ingress interface of the border router. The packets are measured against a traffic profile including a rate and a time over which the rate is measured. However, we do not further address the timing issue in this draft.

The policer causes each packet arriving into the domain to be marked with two levels of drop precedence, which we refer to in the following as green and yellow (in increasing order). The packets conforming to the traffic profile, MUST be marked green (low drop precedence). The excess packets MUST be marked as yellow (high drop precedence). Yellow marked packets MAY be dropped at the ingress router.

The green packets MUST be marked with the DSCP for AFx1. Yellow packets MUST be marked with DSCP for AFx2 or AFx3. X MUST be any value from 1 to N, where N=4 for general use [AF-PHB].

4.2. PHB Configuration

The one-to-any AR traffic aggregate is to be treated using one of the two PHBs in the selected AF PHB class.

Interfaces internal to the domain SHOULD not drop packets marked to receive treatment with AFx1. A node MUST start dropping AFx2 and AFx3 packets before start dropping AFx1. The drop probability and starting point (buffer fill level) of AFx2 and AFx3 MAY be the same.

In the case where the AF class is lightly loaded AFx2 and AFx3 packets SHOULD be transmitted successfully through the node.

At each node, a certain portion of the forwarding resources should be pre-allocated for the AF class. The level of this resource should

not be pre-empted by other PHBs.

A certain number of the N AF classes may be used for this one-to-any AR PDB. Services using the same AF class have only one defined loss probability.

4.3. Admission rule

Since the proposed PDB aims at assuring a certain rate with a given probability, the admission rule for the proposed service needs to be based on probability computations. In this section we propose an admission rule based on the following assumptions:

- All sources are sending at their full assured rate (busy hours) bounded by the policer in the ingress router (see the edge rule). This is the worst case to be taken into account. The traffic aggregate at an ingress router is split towards the egresses of a domain. The splitting can follow different patterns. For example, a traffic aggregate with an assured rate of 6Mbps in a domain with 3 egresses could send with the following patterns: pattern 1) 2 Mbps to each egress; pattern 2) 1 Mbps to egress 1, 2 Mbps to egress 2 and 3 Mbps to egress 3; and so on. The timeframe over which these patterns are considered should be short enough such that the splitting of traffic towards the egresses does not change significantly over that timeframe.

- The probability associated to each of the patterns is known (if it is not known, it can be estimated as discussed below).

- The probability of not meeting the assured rate commitments (i.e. of dropping assured traffic) is given as a quality parameter.

With the above assumptions, admission of a new traffic aggregate can be decided with the following steps:

- In each link determine the probability that the link receives an amount of assured rate traffic larger than the service class' capacity (probability of assured rate congestion in a link). Note that we assume a buffer-less model.

- From the probabilities of assured rate congestion for all the links in the network, determine the probability that a traffic aggregate receives a lower assured rate than the specified (violation probability of the rate assurance).

- Enforce that for all traffic aggregates, the violation probability of the rate assurance is smaller than the probability set as quality parameter.

The main problem in the above assumptions is to determine the probabilities associated to the different patterns. One solution may

be the estimation of these probabilities. For example, the same model of the probability density may be assumed for all traffic aggregates, and it may be weighted by the total traffic rate measured at the corresponding egress.

5. Attributes

Attributes of the PDB include:

- The throughput, defined as the sum of the traffic from one ingress to all egress nodes.

- The probability that no assured rate traffic is dropped (probability of no PDB violation).

6. Parameters

This PDB MUST have the following parameters:

- A policer rate that decides the marking of packets.

In addition to the above, the PDB MAY have optional extra traffic parameters, namely the egress distribution of the traffic, or alternatively the probability of traffic going to a specific egress. These parameters can help to determine the probabilities associated to different patterns for egress distributions of traffic. In case these parameters are not given, these probabilities will have to be estimated by other means (see section 4.3).

Note that traffic parameters specifying the timing relation of the policer rate may be used, but since they are based on a very fine granularity, we do not address them in this draft. Examples of such traffic parameters include a Committed Burst Size (CBS) and an averaging interval (T1).

7. Assumptions

New users without an egress distribution specified will behave as the average of the users already in the network in terms of the egress distribution.

Statistical behavior of the individual user traffic is known in terms of distribution towards egress nodes. If not known we can assume that the distribution is proportional to the utilization or the capacity of the outgoing link at the egress nodes. However, this assumption is only correct if the link capacity of links leaving the DS domain is properly planned.

Furthermore, we assume near static routing or route pinning with mechanisms like using MPLS beneath IP.

8. Example Uses

An example may be a web site that wants to provide its users with high-speed access to its web pages. E.g., a company that sells software/videos/electronic contents via web and is willing to pay in

order to let its users comfortably download the software/videos/electronic contents at high speed. The SLA for this service could be e.g. an aggregated assured average rate of 2 Mbps toward any egress. Note that, because of the nature of this service, the proper support of TCP in such cases is very important, but is out of scope for this document.

Another example could be corporate Internet access services. An enterprise whose business is based on the Internet and is willing to pay in order to provide its workers with high speed Internet access. The sending part of the one-to-any service is covered by this document. The receiving part (any-to-one) is for further study and may result in a different PDB definition.

9. Security Considerations

TBD

10. Open Issues/ TBD

If one-to-any AR PDB(i) uses AFx class in the DS domain and one-to-any AR PDB(j) uses AFy class and the probability of getting the rate of PDB(i) is higher than PDB(j) x MUST be less than y. Does this restriction make sense?

An analytical approach will be used to determine the acceptance region, based on the probabilistic description of the egress distribution inside the network. Simulation at the aggregate flow level will be used to validate the analytical outcome, focusing on the border of the acceptance region.

Are we time dependent? E.g., does the time over which the rate is assured play a role?

Discrete versus continuous egress distribution? Assuming gaussian distribution, where the mean rate on a link is the mean value of the gaussian distribution. How to find the variance?

11. References

[AR-PDB] N. Seddigh, B. Nandy, J. Heinanen, "An Assured Rate Per-Domain Behavior for Differentiated Services", <draft-seddigh-pdb-ar-00.txt>, December 2000.

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

[PDBDEF] K. Nichols, B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their

Specification", <draft-ietf-diffserv-pdb-def-03>,
January 2001.

12. Authors' Addresses

Marcus Brunner, Albert Banchs, Sandra Tartarelli
NEC Europe Ltd.
C&C Research Laboratories
Adenauerplatz 6
D-69115 Heidelberg, Germany
Phone: +49 (0)6221 905110
Fax: +49 (0)6221 9051155
Email: [brunner|tartarelli|banchs]@ccrle.nec.de

Huanxu Pan
NEC Corporation
Network Development Laboratories
1131, Hinode, Abiko, Chiba, 270-1198, JAPAN
Phone: +81-471-85-6737
Fax: +81-471-85-6841
Email: h-pan@cb.jp.nec.com