

Mapping to ATM classes of service for
Differentiated Services Architecture

[draft-ayandeh-diffserv-atm-00.txt](#)

Expiration: April 2000

S. Ayandeh
A. Krishnamurthy
A. Malis
Lucent Technologies

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

The guidelines for PHB specifications contained in the Differentiated Services (DS) Architecture [1] require descriptions of:

- [1](#). How a PHB would map to different link layers
- [2](#). How a PHB would inter-work with non-DS compliant nodes and networks

This draft includes the mapping to ATM classes of service for EF [[2](#)] and AF PHBs [[3](#)].

65

2.0	Introduction	3
3.0	Inter-working issues of DiffServ and ATM	3
4.0	Differentiated Services Requirements for ATM Classes of Service	4
5.0	Recommended Mapping of DiffServ PHBs to ATM Classes of Service	4
5.1	Mapping of EF PHB to CBR Class of Service	4
5.2	Mapping of AF PHB to ABR Class of Service	5
6.0	Example Use of Virtual Circuits	6
7.0	Interactions of TCP and UDP	6
8.0	Security Considerations	7

[2.0](#) Introduction

Multi-service networks form part of the mosaic of existing and emerging data networks. As such there is a need to ensure that characteristics of IP services based on Differentiated Services (DS) architecture are maintained end to end across (intermediate) ATM networks. This is no easy task, as ATM networks standardize a service at the user network interface while allowing for adjustments to a set of well-defined traffic parameters. DS, on the other hand, supplies the supporting traffic conditioning and per-hop behaviors and leaves the service specification to be defined by the service provider.

The goal of this document is to specify the mapping of differentiated services PHBs to existing ATM classes of service. The ATM traffic class, its descriptors, and QoS parameters, required to carry a given PHB are discussed. Circuit aggregation issues are limited to a discussion of virtual path connections.

The motivation for this approach is:

- a) To allow for traffic conditioning functions to occur at frame boundaries while the ATM network delivers a given PHB using existing ATM classes of service. Note that ATM Forum's draft addendum [\[5\]](#), which proposes extensions to ATM service categories and signaling scheme, deals with new mechanisms that are in process of being specified. These mechanisms in effect make an ATM switch DiffServ compliant.
- b) To ensure that the resulting mappings do indeed meet the requirements of the EF and AF PHBs outlined in RFCs 2598 & 2597. For example as we show further on, the approach also recommended by the ATM Forum's draft addendum [\[5\]](#) to their Traffic Management 4.1 Specification [\[4\]](#), of mapping IP services to ATM classes of service may at best approximate the expected behaviors and the resulting IP service.
- c) To simplify the resulting deployments by offering two existing ATM classes of service that would correspond to the EF and AF per hop behaviors. We label this as the "PHB-mapping" approach. This is in contrast to the "service-mapping" approach taken in the addendum to TM 4.1 [\[5\]](#). In the latter approach, services resulting from the AF PHB alone may be mapped to five ATM classes of service: rt-VBR, nrt-VBR, ABR, GFR, and Differentiated UBR.

[3.0](#) DiffServ and ATM inter-working issues

The goal is to meet the requirements of a given differentiated services PHB using a minimum set of resources in an ATM network. Some clarification of the terms used to describe the behavior of the traffic source & QoS parameters in ATM networks is in order. This clarification is intended to highlight the main differences between cell and packet

based metrics currently in use:

Cell rates: Cell rates have to take into account the adaptation layer overheads of ATM such as padding. Adaptation overheads vary with the size of the packets. Therefore conversion of a given bit rate at a packet interface to cell rate has to take into account the statistical nature of the packet length. Some additional bandwidth may have to be allocated in order to account for the statistical nature of the packet length.

Cell Transfer Delay (CTD): This parameter in TM 4.1 is measured as the

65

interval between a pair of cell entry and exit events. CTD does not cover the transmission time at the egress port. In DS measurements the packet transmission time, which varies with packet size, is often included.

Peak-to-peak CDVT: This parameter is measured as the difference $CTD_x - CTD_d$, where subscript x refers to an arbitrary cell and d to a defined reference cell [ITU-T I.356] (injected at idle time). The EF-PHB, for example, uses the absolute value of the difference in nodal CTD for two adjacent packets as a measure of jitter experienced at a node.

Cell Loss Ratio (CLR): CLR represents a lower bound on the loss probability of the corresponding packets. Settings of CLR would require some understanding of the cell loss process and how it relates to the packet loss ratio. For example, partial and early packet discard mechanisms can realize this lower bound.

[4.0](#) Differentiated Services Requirements for ATM Classes of Service
Currently defined DiffServ architecture places three basic requirements on the ATM network: marking of packet drop precedence, as well as, minimum cell rate control and active queue management.

[1](#). Marking of Packet Drop Precedence: ATM offers a cell loss priority (CLP=0 or 1) mechanism. Cells resulting from in profile packets are marked with CLP=0, while cells resulting from out of profile packets are marked with CLP=1. Marking of cells should therefore be applied to entire frames and occur at frame boundaries. Indiscriminate marking of cells would lead to unacceptable throughput behavior in the face of congestion in the cell network [\[6\]](#).

[2](#). Minimum Cell Rate [cell/s]: Needs to be configured for both the EF and AF PHBs. For EF it ensures that the aggregate has a well-defined minimum departure rate over a time interval equal to or greater than the time it takes to send an output link MTU sized packet at the configured rate of the EF-PHB. This together with policing action at the edge leads to a low loss, delay, and jitter per-hop behavior.

3. Active Queue Management: The assured forwarding PHB and hence the resulting services, require the properties of a RED like algorithm for active queue management [7]. ATM services can offer cell discard for CLP=1 cells. However, indiscriminate cell drops are not of any use in support of AF PHB. Both partial and early packet discard mechanisms in ATM lead to a tail packet drop behavior. Tail drop is undesirable for its adverse effects on TCP flows. ATM traffic classes [4] are being enhanced through local policy to offer active queue management. However, any solution to this problem is further aggravated by virtual path connections where individual circuit and frame delineation are not visible by definition.

5.0 Recommended Mapping of DiffServ PHBs to ATM Classes of Service
The requirements outlined in the previous section are met by the following mappings of differentiated services PHBs to existing ATM classes of service.

5.1 Mapping of EF-PHB to the rt-VBR Class of Service
Source Traffic Descriptor (VBR.1)

65

Peak Cell Rate (PCR) (CLP=0+1) = line rate of the connection with its inverse T_{0+1}

Sustained Cell Rate (SCR) = configured rate of EF PHB with its inverse T_s

Maximum Burst Size (MBS) = max-PDU-size

Two conformance definitions apply GCRA(T_{0+1} , CDVT) to regulate the peak and GCRA(T_{s0} , $BT_0 + CDVT$) to regulate the rate of the connection averaged over the life of the connection [4]. Non conforming cells are dropped.

QoS Parameters

CLR = Cell loss is not expected along the ATM path, however loss may occur at the ingress for non-conforming packets

MaxCTD = Is the transfer delay of the ATM network

Peak-to-peak CDV = a derivative of MaxCTD is equal to MaxCTD-(fixed propagation, transmission, and switching delays)

Note that the network through the traffic descriptor provisions adequate resources to accommodate cells bursting at line rate at the configured rate of the EF-PHB. By using an MBS equal to the max-PDU-size, the network over provisions bandwidth. However since EF-PHB is a policed service, the excess bandwidth would be available for use by other connections.

Mapping of EF-PHB to the CBR class of service would introduce shaping induced delays at egress when re-assembling a packet. Low delay

requirement of EF-PHB may therefore not be well served with CBR class of traffic.

[5.2](#) Mapping of AF PHB to ABR Class of Service

Source Traffic Descriptor (ABR)

PCR (CLP=0+1) = minimum line rate along the path of the VC with its inverse T_{0+1}

MCR = minimum bandwidth allocated to an AF class with its inverse T_m0

The conformance definitions is GCRA(T_{m0} , CDVT), or based on ideal transmission time of CLP=0 cells resulting from the "Allowed Cell Rate" in response to the feedback mechanism. All data cells are transmitted with CLP=0. Congestion feedback is supplied as an explicit rate, congestion indication, or no increase parameters contained in resource management cells.

QoS Parameters

CLR = no value needs to be specified. The target value is however intended for conforming connections and is network specific.

The use of ABR traffic class is recommended as it allows all of the requirements of the AF-PHB to be met. Note that VBR, GFR, and relative UBR only partially meet the packet drop precedence requirement, and do not offer active queue management. Furthermore, these traffic classes do not support virtual path circuits for reasons already discussed in section-3.

65

With ABR, the minimum throughput and buffer requirements of the AF PHB can be provisioned. With a minimum amount of cell loss, congestion in

the cell network is pushed back to the frame boundaries where active queue management is applied. Any mechanisms to implement back pressure at the frame boundary are implementation specific.

Active queue management is applied to the full range of the packet drop precedence markings i.e. AFx1, AFx2, and AFx3. This is in contrast to the three to two levels mapping which would occur with the use of Cell Loss Priority (CLP) in the service-mapping approach.

The requirements of the AF PHB are therefore fully met. Given that only a few ABR connections would be carrying large aggregates of AF traffic, any possible concern with the volume of resource management cells does not arise. For list of numerous references regarding IP traffic and ABR see references [[8](#), [9](#), [10](#)].

[6.0](#) Example Use of Virtual Circuits

Please note that the code points are only mentioned for illustrative

purpose. The idea is to map PHBs to ATM classes of service and not DS code points.

Diffserv code point		VC-type	
101110	EF	rt-VBR	VCb
001010	AF11	ABR	VCc CLP=0
001100	AF12		CLP=0
001110	AF13		CLP=0
010010	AF21	ABR	VCd CLP=0
010100	AF22		CLP=0
010110	AF23		CLP=0
011010	AF31	ABR	Vce CLP=0
011100	AF32		CLP=0
011110	AF33		CLP=0
100010	AF41	ABR	VCf CLP=0
100100	AF42		CLP=0
100110	AF43		CLP=0
000000	BE	UBR	VCg

7.0 Interactions of TCP and UDP

There are three alternatives available to handle the interactions of TCP and UDP traffic [[11](#), [12](#)]:

a) Integrated treatment with two drop precedence:

TCP/UDP In Profile -> AFx1

TCP/UDP Out Profile -> AFx2

In this scenario both TCP and UDP would achieve their committed rates. In the face of persistent congestion UDP would consume the entire excess bandwidth.

65

b) Integrated treatment with three drop precedence:

This scenario can not be supported by the "service-mapping" approach as ATM is limited to two drop precedence levels. However the "PHB-mapping"

approach can support this scenario:

TCP/UDP In Profile -> AFx1

TCP Out Profile -> AFx2

UDP Out Profile -> AFx3

In this scenario both TCP and UDP would achieve their committed rates. In the face of persistent congestion, TCP and UDP would share the excess

bandwidth. This assumes that some excess bandwidth has been provisioned. WRED algorithms that maintain per drop level packet counts further aid the fair allocation of the excess bandwidth.

c) Separate treatment (with two drop precedence)

TCP and UDP may each be treated as a separate class, given the inadequate sharing of excess bandwidth in scenario "a" above. For example:

TCP in profile -> AFx1

TCP out profile -> AFx2

UDP in profile -> AFy1

UDP out profile -> AFy2

Use of ABR to support the AF-PHB enables scenario "b" the "integrated treatment with three drop precedence". This saves on the number of AF classes in use and hence the number of virtual circuits required.

8.0 Security Considerations

As with any other provisioned services, the ATM network must make use of its capabilities for call admission control and police against sources attempting to utilize more network resources than their service contract allows. In addition, when the ABR service class is used, the network provides flow control feedback to sources. This allows these sources to avoid data loss on the ATM network if congestion occurs. Congestion may be due to denial of service attacks from other sources or temporary network outages.

9.0 References

[1] Blake, S., Black, D., Carlson, M., Davis, E., Wang, W., Weiss, W., "An Architecture for Differentiated Services", [RFC 2475](#), December 1998.

[2] Jacobson, V., Nichols, K., Poduri, K., "An Expedited Forwarding PHB", [RFC 2598](#), June 1999.

[3] Heinanen, J., Baker, F., Weiss, W., Wroclawski, J, "Assured Forwarding PHB Group", [RFC 2597](#), June 1999.

[4] ATM Forum Traffic Management 4.1 Specification, ATM Forum/af-tm-0121.000, March 1999.

[5] Addendum to TM 4.1: Enhancements to Support IP Differentiated

[6] Romanow, A., Floyd, S., "Dynamics of TCP traffic over ATM networks", Proceedings of SIGCOMM'94, September 94.

[7] Braden, B., Clark, D., Crowcroft, J. , Davie, B. , Deering, S. , Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., Zhang, L., Recommendations on Queue Management and Congestion Avoidance in the Internet, April 1998. Available as [RFC 2309](#) (text) as an Informational RFC

[8] Fahmy, S., Jain, R., Rabie, S., Goyal, R., Vandalore, B., "Quality of Service for Internet Traffic over ATM Service Categories," Journal of Computer Communications special issue on enterprise networks, 1999

[9] Vandalore, B., Kalyanaraman, S., Jain, R., Goyal, R., Fahmy, S., "Simulation Study of World Wide Web traffic over the ATM ABR Service," Proceedings of SPIE Symposium on Voice, Video and Data Communications, Vol. 3530, Conference on Performance and Control of Network Systems II, Boston, MA, November 1998, pp. 415-422

[10] Pazos, C. M. D., Signore, V. A., Cavendish Jr., D., Gerla, M., "Performance of TCP over ATM for Various ABR Control Policies," Proceedings of ICCCN'96, Rockville, MD. October, 1996.

[11] Goyal, M., Durresi, A., Jain, R., Chunlei, L., "Effect of Number of Drop Precedence in Assured Forwarding", [draft-goyal-diffserv-dpstdy-02](#), July 1999

[12] Elloumi, O., Cnodder, S., Pauwels, K., "Usefulness of three drop precedence in Assured Forwarding service", [draft-elloumi-diffserv-threestwo-00.txt](#), July 1999

[10.0](#) Authors' Addresses

Siamack Ayandeh
Lucent Technologies
[1](#) Robbins Road,
Westford, MA, 01886
(978) 952 7866
sayandeh@ascend.com

Anand Krishnamurthy
Lucent Technologies
[1](#) Robbins Road,
Westford, MA, 01886
(978) 952 1448
ak26@lucent.com

Andrew Malis
Lucent Technologies
[1](#) Robbins Road,

Westford, MA, 01886
(978) 952 7414
amalis@lucent.com