The ATM Forum
Technical Committee

Addendum to TM 4.1: Differentiated UBR

af-tm-0149.000

July, 2000

The ATM Forum
Worldwide Headquarters
2570 West El Camino Real, Suite 304
Mountain View, CA 94040-1313
Tel: +1-650-949-6700
Fax: +1-650-949-6705

Table of Contents

1. Introduction

The UBR service category is often utilized by elastic data networking applications, primarily because it imposes few requirements on the end equipment and enables a higher degree of connectivity (i.e., more connections can be configured on a given access interface) than do services which explicitly reserve capacity per connection.

Such a configuration offers no means by which the service quality of individual flows (virtual channel or virtual path connections) may be differentiated. This may not be an issue if the higher layer intends to offer only best effort service. It has, however, become increasingly common for technologies which traditionally offered only best effort service (e.g., Ethernet, IP) to incorporate service differentiation mechanisms. In order to optimize their support over ATM, it is desirable that the UBR service category be extended with similar capabilities.

This specification defines a mechanism by which a UBR connection may be associated with one of a set of network-specific *Behavior Classes* (see section 2.5). Using this mechanism, the adaptation function may assign traffic to a UBR connection, based in part on a mapping of the higher-layer service classification to the Behavior Class with which that UBR connection is associated. The ATM network may differentiate the treatment of UBR connections based on the Behavior Class with which each such connection is associated, for example by governing access to queuing and scheduling resources. By coordinating the adaptation function at the edge of the ATM network with the behaviors implemented within the network, an operator can enable consistent service differentiation end to end.

UBR connections which do not indicate a Behavior Class are associated with a network-specific default behavior. Networks which do not support the specification of Behavior Class ignore this attribute and handle the connection as if it had not been specified. The enhancements defined herein are optional and interoperable with TM4.1 compliant UBR implementations.

This specification uses the term ***Differentiated UBR*** to refer to the UBR service, augmented by the specification of a Behavior Class.

1.1 Document Structure

Section 2 defines extensions and modifications required to TM 4.1, in order to enable support of the ***Differentiated UBR*** capability as defined in this specification.

Section 3 provides details of documents referenced herein.

Section 4 defines acronyms used in this specification.

Appendix A discusses the IETF's Differentiated Services (diffserv) architecture, and describes how the mechanisms defined in this specification may be used to support it.

Appendix B discusses the traffic prioritization capability defined in the IEEE's 802.1D standard, and describes how the mechanisms defined in this specification may be used to support it.

Appendix C discusses use of Differentiated UBR to support other service policies.

2. Modifications to TM 4.1

2.1 Section 2.1.4, Unspecified Bit Rate (UBR) Service Category Definition

Modify this section as follows:

> The Unspecified Bit Rate (UBR) service category is intended for non-real-time applications, i.e., those not requiring tightly constrained delay and delay variation. Examples of such applications are traditional computer communications applications, such as file transfer and email.

UBR service does not specify traffic related service guarantees. No numerical commitments are made with respect to the CLR experienced by a UBR connection, or as to the CTD experienced by cells on the connection. A network may or may not apply PCR to the CAC and UPC functions. In the case where the network does not enforce PCR, the value of PCR is informational only. When PCR is not enforced it is still useful to have PCR negotiated, since this may allow the source to discover the smallest bandwidth limitation along the path of the connection. Congestion control for UBR may be performed at a higher layer on an end-to-end basis.

A UBR connection may optionally be associated with a Behavior Class [see section 5.11] chosen from a set supported by the network. The network may in turn use the indicated Behavior Class in determining the treatment to be applied to that connection. Neither the definition of a set of Behavior Classes, nor the corresponding network behaviors, are within the scope of this specification. The term Differentiated UBR is used in this specification to refer to the UBR service category augmented by the specification of a Behavior Class.

The UBR service is indicated by use of the Best Effort Indicator in the ATM User Cell Rate Information Element.

2.2 *Section 2.2, "ATM Service Category Parameters and Attributes"*

Modify Table 2-1 as follows:

	ATM Layer Service Category					
Attribute	CBR	rt-VBR	nrt-VBR	UBR	ABR	GFR
Traffic Parameters[4]						
PCR and CDVT[5]	Specified			Specified[2]	Specified[3]	Specified
SCR, MBS, CDVT[5]	n/a	Specified		n/a		
MCR	n/a				Specified	n/a
MCR, MBS, MFS, CDVT[5]	n/a					Specified
QoS Parameters[4]						
Peak-to-peak CDV	Specified		Unspecified			
MaxCTD	Specified		Unspecified			
CLR	Specified			Unspecified	See note 1	See note 7
Congestion Control[8]						
Feedback	Unspecified				Specified[6]	Unspecified
Other Attributes						
BCS	Unspecified			Optional	Unspecified	

Table 2-1: ATM Service Category Attributes

Notes:

1. CLR is low for sources that adjust cell flow in response to control information. Whether a quantitative value for CLR is specified is network specific.
2. Might not be subject to CAC and UPC procedures.
3. Represents the maximum rate at which the ABR source may ever send. The actual rate is subject to the control information.
4. These parameters are either explicitly or implicitly specified for PVCs or SVCs.

5. CDVT refers to the Cell Delay Variation Tolerance (see Section 4.4.1). CDVT is not signaled. In general, CDVT need not have a unique value for a connection. Different values may apply at each interface along the path of a connection.
6. See Section 2.4.
7. CLR is low for frames that are eligible for the service guarantee. Whether a quantitative value for CLR is specified is network specific.
8. See Section 5 for details of ABR feedback and other congestion control mechanisms.

2.3 Section 2.3.1, "Nature of Service Guarantees"

Modify the second paragraph to read as follows:

> The UBR service category offers no traffic related service commitments. No numeric commitment is made as to the cell loss ratio experienced by a connection, or as to the cell transfer delay experienced by cells on the connection. Fairness among connections cannot be assumed, although local policy in some network elements may have this effect. The treatment accorded to an individual connection may also be influenced by the Behavior Class value associated with that connection, when such a value is specified.

2.4 Section 2.3.2 "Mechanisms"

Modify the second paragraph to read as follows:

> The UBR service category is inherently open loop. UBR is not subject to a specific traffic contract but may be subject to a local policy in individual switches and end-systems. The treatment accorded to an individual connection may be influenced by the Behavior Class value associated with that connection, when such a value is specified.

2.5 New section 5.11, "UBR Behavior Class"

Add the following to text to section 5 of TM 4.1.

5.11 UBR Behavior Class

The service provided to UBR connections may be differentiated by associating each such connection with one of a set of network-specific groups, referred to here as *behavior classes*. In such an association, the behavior class is indicated via the Behavior Class Selector (BCS) parameter. UBR connections for which no behavior class is indicated are associated with a network-specific default behavior. The use of BCS with ATM service categories other than UBR is for further study.

Network resources (e.g., queuing and scheduling resources) may also be associated with a behavior class, as a means of enabling service differentiation amongst behavior classes.

The set of behavior classes, and the corresponding Behavior Class Selector values supported within a given network, are not subject to specification.

This capability is applicable to both VCCs and VPCs, and to all connection topologies (e.g., point-to-point, point-to-multipoint). The behavior class of a UBR connection need not be the same in both directions. It is intended that the behavior class of a UBR connection be consistent at all queuing points[1] in a given direction. It is recommended that BCS values be translated where necessary, e.g., at boundaries between administrative domains, in order to ensure consistent treatment end to end. If such a translation is not

---

[1] A queuing point is a point of resource contention within a network element, where cells may potentially be delayed or lost. A switch may contain multiple queuing points.

performed between devices which assign different meaning to a given BCS value, end to end service may be inconsistent.

Support of the Behavior Class capability is optional. In networks or network elements which do not support specification of a Behavior Class, the indication of Behavior Class is ignored. Connections which do not specify a Behavior Class are associated with a network-specific default behavior. Networks which do not support the specification of Behavior Class ignore this attribute and handle the connection as if it had not been specified.

The BCS parameter differs from the parameters associated with other ATM service categories in the sense that its impact on the end-to-end service is intentionally not specified, and is therefore implementation and/or network specific. The effect of a particular value of the BCS parameter may be different in different networks.

Modification of the association between a UBR connection and a behavior class, after the connection is in the active state, is for further study.

3 Abbreviations

The following abbreviations are used in this specification.

ARP	Address Resolution Protocol
BCS	Behavior Class Selector
CPCS	Common Part Convergence Sublayer
DS	Differentiated Services
DSCP	Differentiated Services Code Point
FDDI	Fiber Distributed Data Interface
MAC	Media Access Control
PHB	Per Hop Behavior
SAR	Segmentation And Reassembly

4 References

[1] IETF RFC 2474, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", Internet Engineering Task Force, December 1998.

[2] IETF RFC 2475, "An Architecture for Differentiated Services", Internet Engineering Task Force, December 1998.

[3] ISO/IEC 15802-3 [IEEE Standard 802.1D], "Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Common specifications - Part 3: Media Access Control (MAC) Bridges", 1998.

[4] ITU-T Recommendation I.363.5, "B-ISDN ATM Adaptation Layer specification: Type 5 AAL"

[5] IETF RFC 2225 "Classical IP and ARP over ATM", Internet Engineering Task Force, April 1998

[6] The ATM Forum, "LAN Emulation over ATM Version 1.0", af-lane-0021.000, January 1995

[7] The ATM Forum, "Multi-Protocol over ATM Version 1.0", af-mpoa-0087.000, July 1997

Appendix A. Support of the IETF Differentiated Services Model

This appendix describes how the Differentiated UBR mechanism may be used to support the IP Differentiated Services (diffserv) model [2].

A.1. Overview of the Differentiated Services model

The diffserv architecture enables IP networks to support QoS differentiation between aggregated flows [2]. The diffserv model is based on traffic classification and conditioning (*e.g.*, policing and marking) at the network edge, and a limited set of per hop behaviors (PHBs) implemented at queuing points within the network. PHBs are externally observable forwarding behaviors applied at a diffserv compliant node to an aggregate flow. Forwarding behaviors include the differential treatment an individual IP packet receives. The differentiated services code point (DSCP) is used to select the PHB that a packet experiences at a given queuing point[1]. The DSCP is encoded in the DS field of the IP header.

Diffserv and ATM differ with respect to the nature of their respective QoS objectives. ATM QoS objectives are always absolute; *i.e.*, QoS measures must be specifically defined in terms of loss and delay, and the assurances (if any) offered to a connection are independent of those offered to other connections. Diffserv can be used to provide services with absolute QoS objectives, and services with relative service objectives. Services with absolute QoS objectives are similar to the ATM service categories in that their QoS measures must be specifically defined. Services with relative service objectives have the property that the service level provided to one connection may depend on that provided to other connections.

A.2. Use of Differentiated UBR to support Diffserv

This section describes one way in which diffserv may be supported over ATM, using Differentiated UBR.

The general architecture presumes a set of IP devices attached to a common ATM network. Various mechanisms exist to establish their interconnection, including Classical IP and ARP over ATM [5], LAN Emulation (LANE) [6], MPOA [7] and manual configuration.

To support diffserv, multiple UBR virtual channel connections are established between each pair of ATM-attached IP devices. A behavior class representative of the desired treatment within the ATM network is assigned to each UBR connection. Knowledge of this assignment is retained by the IP layer of the devices at the end points of the connection. Finally, a mapping table is instantiated which translates Differentiated Services Code Points (DSCPs) to behavior classes and loss priority (i.e., drop precedence).

To forward a given packet, the ATM-attached IP device first determines the device to which the packet should be sent. The DSCP in the packet header is used to determine the corresponding behavior class and loss priority. In order to honor packet ordering constraints, the translation should be from DSCP to PHB Scheduling Class (PSC) and then from PSC to BCS. A PSC is a set of PHBs which share the same ordering constraints, while each PHB is specified by the DSCP. Finally, the combination of “next hop” device and behavior class is used to select a UBR connection over which to forward the packet. The loss priority is communicated to the ATM adaptation layer via the *CPCS-Loss Priority* parameter of the *CPCS-UNITDATA.Invoke* primitive [4], and coded into the CLP bit of the headers of the corresponding cells by the Segmentation and Reassembly (SAR) Sublayer.

Queuing points within the ATM network treat a cell according to the loss priority indicated in its header and the behavior class of the UBR connection on which it arrives.

Appendix B. Support of IEEE 802.1D User Priorities

This appendix describes how the Differentiated UBR mechanism may be used to support the User Priorities as specified in IEEE 802.1D "Media Access Control (MAC) Bridges" [3].

B.1. Overview of the IEEE 802.1D User Priority Mechanism

IEEE 802.1D supports packet classification and service differentiation via eight user priority levels. User priorities are encoded in three bits. Some LANs convey the user priority as an intrinsic part of their MAC frame (*e.g.*, token ring, FDDI), while others encode the user priority in a tag header inserted into the MAC frame (*e.g.*, Ethernet). User priorities are mapped to separate traffic classes within an 802.1D-compliant bridge or host.

Further information can be obtained by referring to Annex H of ISO/IEC 15802-3 ("Design Considerations for Traffic Class Expediting and Dynamic Multicast Filtering")[3].

B.2. Use of Differentiated UBR to support 802.1D User Priorities

MAC bridging over ATM is typically implemented by meshing several MAC layer bridges using ATM virtual connections. Each virtual connection is modeled as a logical interface. Some means, either explicit distribution via the spanning tree protocol or passive monitoring of source addresses in frame headers, is used to form an association of MAC address with logical interface. The outgoing interface through which to forward a given frame is determined based on a lookup of the destination address in its MAC layer header in the resulting forwarding table.

To implement 802.1D user priorities within this framework, multiple UBR virtual connections are provisioned between pairs of MAC layer bridges. A BCS value representative of the desired treatment within the ATM network, is assigned to each UBR connection. Finally, a mapping table is instantiated which translates User Priority values to BCS values.

To forward a given frame, the ATM-attached MAC layer bridge first performs a destination address lookup. This will in general return a set of virtual connections. The User Priority in the MAC layer header is then used in conjunction with the User_Priority-to-BCS mapping table, to determine the corresponding BCS. Finally, the BCS is used to select a UBR connection from amongst the set returned by the address lookup.

At queuing points within the network, the treatment to be provided individual frames is determined from the BCS value associated with the UBR virtual connection on which the corresponding cells arrive.

Appendix C. The use of BCS for other service policies

The usage of BCS in differentiated UBR is not restricted to IETF Diffserv as in Appendix A or IEEE 802.1D user priority as in Appendix B. Differentiation of UBR connections based on another service policy is also possible. The BCS mapping for each connection can in such cases be determined without examining the IP or MAC header. For example BCS may be determined from other information, such as port configuration, physical / logical interface ID or any other higher layer parameters. Specification of such BCS mapping policy is beyond the scope of this document.

End of Document