

Terminology for ATM ABR Benchmarking

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This memo discusses and defines terms associated with performance benchmarking tests and the results of these tests in the context of Asynchronous Transfer Mode (ATM) based switching devices supporting ABR (Available Bit Rate). The terms defined in this memo will be used in addition to terms defined in RFCs 1242, 2285, and 2544 and 2761. This memo is a product of the Benchmarking Methodology Working Group (BMWG) of the Internet Engineering Task Force (IETF).

1. Introduction

This document provides terminology for benchmarking ATM based switching devices supporting ABR. It extends terminology already defined for benchmarking network interconnect devices in RFC's 1242, 2285, and 2544 and 2761. Although some of the definitions in this memo may be applicable to a broader group of network interconnect devices, the primary focus of the terminology in this memo is on ATM ABR.

This memo contains two major sections: Background and Definitions. The background section provides the reader with an overview of the technology and IETF formalisms. The definitions section is split into two sub-sections. The formal definitions sub-section is provided as a courtesy to the reader. The measurement definitions sub-section contains performance metrics with inherent units.

This document assumes that necessary services are available and active. For example, IP connectivity requires SSCOP connectivity between signaling entities. Further, it is assumed that the SUT has

the ability to configure ATM addresses (via hard coded addresses, ILMI or PNNI neighbor discovery), has the ability to run SSCOP, and has the ability to perform signaled call setups (via UNI or PNNI signaling). Finally, this document presents only the terminology associated with benchmarking IP performance over ATM; therefore, it does not represent a total compilation of ATM test terminology.

The BMWG produces two major classes of documents: Benchmarking Terminology documents and Benchmarking Methodology documents. The Terminology documents present the benchmarks and other related terms. The Methodology documents define the procedures required to collect the benchmarks cited in the corresponding Terminology documents.

2. Existing Definitions.

RFC 1242, "Benchmarking Terminology for Network Interconnect Devices", should be consulted before attempting to make use of this document. RFC 2544, "Benchmarking Methodology for Network Interconnect Devices", contains discussions of a number of terms relevant to the benchmarking of switching devices and should be consulted. RFC 2285, "Benchmarking Terminology for LAN Switching Devices", contains a number of terms pertaining to traffic distributions and datagram interarrival. RFC 2761, "Terminology for ATM Benchmarking", contains a number of terms pertaining to traffic management [TM4.0, TM4.1]. Many of the metrics defined in RFC 2761 (e.g., CDV, CER, CLR, CMR, and CTD) also apply to ABR performance benchmarking. These metrics will not be redefined in this document. For the sake of clarity and continuity, this RFC adopts the template for definitions set out in Section 2 of RFC 1242.

II. Definitions

The definitions presented in this section have been divided into two groups. The first group is formal definitions, which are required in the definitions of the performance metrics but are not themselves strictly metrics. These definitions are subsumed from other work done in other working groups both inside and outside the IETF. They are provided as a courtesy to the reader.

1. Formal Definitions

1.1. Definition Format (from RFC 1242)

Term to be defined.

Definition: The specific definition for the term.

Discussion: A brief discussion of the term, its application and any restrictions on measurement procedures.

Specification: The working group and document in which the terms are specified and are listed in the references section.

1.2. Related Definitions.

1.2.1. Allowed Cell Rate (ACR)

Definition: An ABR service parameter, ACR is the current rate (cells/second) at which a source is allowed to send.

Discussion: For ABR traffic, ACR constitutes the actual data throughput for a particular VC. The time change of this value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.2. ACR Decrease Time Factor (ADTF)

Definition: This is the time permitted between sending RM-cells before the rate is decreased to ICR (Initial Cell Rate). The time units are .01 to 10.23 seconds with a granularity of 10 ms.

Discussion: For ABR traffic, ADTF constitutes the time rate of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.3. Additive Increase Rate (AIR)

Definition: An ABR service parameter, AIR controls the rate at which the cell transmission rate increases. It is signaled as AIRF, where

$$\text{AIRF} = \text{AIR} * \text{Nrm} / \text{PCR}.$$

Discussion: For ABR traffic, AIR effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.4. Additive Increase Rate Factor (AIRF)

Definition: Refer to AIR.

Discussion: Refer to AIR.

Specification: AF-TM4.0

1.2.5. Available Bit Rate (ABR)

Definition: ABR is an ATM layer service category for which the limiting ATM layer transfer characteristics provided by the network may change subsequent to connection establishment. A flow control mechanism is specified which supports several types of feedback to control the source rate in response to changing ATM layer transfer characteristics.

Discussion: It is expected that an end-system that adapts its traffic in accordance with the feedback will experience a low cell loss ratio and obtain a fair share of the available bandwidth according to a network specific allocation policy. Cell delay variation is not controlled in this service, although admitted cells are not delayed unnecessarily.

Specification: AF-TM4.1

1.2.6. Available Bit Rate (ABR) Compliance (Conformance)

Definition: ABR connection conformance refers to the behavior specified for ABR destination and switches, but allows for delays between the source and the UNI [UNI3.1, UNI4.0], which may perturb the traffic flow.

Discussion: The cells on an ABR connection applies to CLP=0 cells, which are tested upon arrival. At the arrival point, each cell is identified as conforming or non-conforming. The minimal conformance definition for ABR is $GCRA((1/PCR), t_1)$, where PCR is defined for CLP=0 flow.

Specification: AF-TM4.1

1.2.7. BN

Definition: The BN bit in the RM-cell indicated whether the RM-cell is a BECN cell or not.

Discussion: If BN=0, the RM cells were generated by the source. If BN=1, the RM cells were generated by the destination or a switch.

Specification: AF-TM4.1

1.2.8. CCR

Definition: The CCR field in the RM-cell is set by the source to its current ACR. CCR is formatted as a rate.

Discussion: For BECN cells, CCR=0.

Specification: AF-TM4.1

1.2.9. Cell Blocks (CB)

Definition: Cell blocks are a sequence of N cells transmitted consecutively on a given connection.

Discussion: A cell block will normally correspond to the number of information cells transmitted between successive OAM cells.

Specification: AF-TM4.1

1.2.10. Congestion Indication (CI)

Definition: The CI bit in the RM-cell allows a network element to indicate that there is congestion in the network.

Discussion: When the source receives a backward RM-cell with CI=1, ACR is decreased. When the destination turns around a forward RM-cell, the CI is set to 1 to indicate that the previously received data cell had the EFCI state set.

Specification: AF-TM4.1

1.2.11. Cutoff Decrease Factor (CDF)

Definition: CDF controls the decrease in ACR (Allowed Cell Rate) associated with CRM (missing RM cell count).

Discussion: For ABR traffic, CDF effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.12. DIR

Definition: The DIR bit in the RM-cell indicates which direction of data flow is associated with the RM-cell. DIR is changed from 0 to 1 when an RM-cell is turned around at the destination.

Discussion: A forward RM-cell is indicated by DIR=0 and is associated with data cells flowing in the same direction. A backward RM-cell is indicated by DIR=1 and is associated with data cells flowing in the opposite direction.

Specification: AF-TM4.1

1.2.13. Explicit Rate (ER)

Definition: The ER field in the RM-cell is used to limit the source ACR to a specific value. For each RM-cell, ER is set by the source to a requested rate (e.g., PCR). It may be reduced by any network element in the path to a value that the element can sustain. ER is formatted as a rate.

Discussion: None.

Specification: AF-TM4.1

1.2.14. Feedback

Definition: Information carried in the backward RM-cells provided by the network elements and/or the destination back to the source.

Discussion: Feedback may include information in the ER field, or the CI or NI bits of each backward RM-cell.

Specification: AF-TM4.1

1.2.15. Ideal Transmission Time (ITT)

Definition: The transmission time for ABR CLP=0 cells, if the difference between itself and the transmission time for the previous CLP=0 cell on the connection is greater than or equal to the minimum:
a) the inverse of the ACR in effect immediately after the transmission time of the first of the two cells
b) the inverse of the ACR in effect immediately before the transmission time of the second of the two cells.

Discussion: The transmission time for the first cell on the connection is automatically an ITT.

Specification: AF-TM4.1

1.2.16. Initial Cell Rate (ICR)

Definition: An ABR service parameter, in cells/sec, that is the rate at which a source should send initially and after an idle period.

Discussion: none.

Specification: AF-TM4.0

1.2.17. In-Rate Cells

Definition: In-Rate ABR cells are sent with CLP=0.

Discussion: ABR RM-cells shall be sent with CLP=0 except in certain circumstances, See Out-of-Rate Cells. All other ABR cells shall be sent with CLP=0.

Specification: AF-TM4.1

1.2.18. Minimum Cell Rate (MCR)

Definition: An ABR service traffic descriptor, in cells/sec, that is the rate at which the source is always allowed to send.

Discussion: MCR may be set to zero. The bandwidth available from the network may vary, but shall not become less than MCR.

Specification: AF-TM4.1

1.2.19. Mrm

Definition: An ABR service parameter that controls allocation of bandwidth between forward W-cells, backward RM-cells, and data cells.

Discussion: none.

Specification: AF-TM4.0

1.2.20. No Increase (NI)

Definition: The NI bit in the RM-cell is used to prevent a source from increasing its ACR. NI does not require any decrease in value.

Discussion: None.

Specification: AF-TM4.0

1.2.21. Nrm

Definition: An ABR service parameter, Nrm is the maximum number of cells a source may send for each forward RM-cell.

Discussion: none.

Specification: AF-TM4.0

1.2.22. Out-of-Rate Cells

Definition: Out-of-Rate ABR cells are sent with CLP=1.

Discussion: This may be used to enable a rate increase for a connection that has an ACR=0. The source would generate out-of-rate cells to probe the network to learn when it may increase its rate.

Specification: AF-TM4.1

1.2.23. Rate Decrease Factor (RDF)

Definition: An ABR service parameter, RDF controls the decrease in the cell transmission rate. RDF is a power of 2 from 1/32,768 to 1.

Discussion: For ABR traffic, RDF effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.24. Rate Increase Factor (RIF)

Definition: This controls the amount by which the cell transmission rate may increase upon receipt of a RM-cell. The additive increase rate $AIR = PCR * RIF$. RIF is a power of 2, ranging from 1/32,768 to 1.

Discussion: For ABR traffic, RIF effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.25. Resource Management (RM) Cells

Definition: RM cells are used to convey network status (available bandwidth, congestion levels) and request peak cell rates for ATM blocks. The RM cell has the following format:

Header: 5 bytes, same as the ATM cell header Protocol ID: 3 bytes, protocol ID value is 1 for ABR service Function specific field: 45 bytes, data required for the specific protocol (See DIR, BN, CI, NI, ER, CCR, and MCR for field information.) Rsvd: 6 bytes, reserved for future specification EDC: 10 bytes, CRC-10 error detection code computed over the cell payload .ti6 (except the CRC-10 field) and used to check for data corruption

Discussion: RM information can exist at the VP and/or VC level. VP level cells are identified with a VCI value of 6. VC level cells are identified with a PT of 6. See DIR, BN, CI, NI, ER, CCR, and MCR for additional protocol field information.

Specification: AF-TM4.0

1.2.26. Severely Errored Cell Block (SECB)

Definition: A severely cell block outcome occurs when more than M errored cells, lost cells, or misinserted cell outcomes are observed in a received cell block.

Discussion: none.

Specification: AF-TM4.1

1.2.27. Tagged Cell Rate (TCR)

Definition: An ABR service parameter, TCR limits the rate at which a source may send out-of-rate forward RM-cells. TCR is a constant fixed at 10 cells/second.

Discussion: none.

Specification: AF-TM4.0

1.2.28. TDF

Definition: An ABR service parameter, TDF controls the decrease in ACR associated with TOF. TDF is signaled as TDFF, where $TDF = TDFF / RDF$ times the smallest power of 2 greater or equal to PCR. TDF is in units of 1/seconds.

Discussion: For ABR traffic, TDF effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.29. TDFF

Definition: Refer to TDF. TDFF is either zero or a power of two in the range 1/64 to 1 in units of 1 /cells.

Discussion: Refer to TDF.

Specification: AF-TM4.0

1.2.30. Time Out Factor (TOF)

Definition: An ABR service parameter, TOF controls the maximum time permitted between sending forward RM-cells before a rate decrease is required. It is signaled as TOFF where $TOF = TOFF + 1$. TOFF is a power of 2 in the range: 1/8 to 4,096.

Discussion: For ABR traffic, TOF effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.31. Time Out Factor (TOFF)

Definition: Refer to TOF.

Discussion: none.

Specification: AF-TM4.0

1.2.32. Trm

Definition: An ABR service parameter that provides an upper bound on the time between forward RM-cells for an active source. It is 100 times a power of two with a range of $100 \cdot 2^{-7}$ to $100 \cdot 2^0$

Discussion: For ABR traffic, Trm effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.33. Virtual Source/Virtual Destination (VSND)

Definition: An ABR connection may be divided into two or more separately controlled ABR segments. Each ABR control segment, except the first, is sourced by a virtual source. A virtual source implements the behavior of an ABR source endpoint. Backward RM-cells

received by a virtual source are removed from the connection. Each ABR control segment, except the last, is terminated by a virtual destination. A virtual destination assumes the behavior of an ABR destination endpoint. Forward RM-cells received by a virtual destination are turned around and not forwarded to the next segment of the connection.

Discussion: none.

Specification: AF-TM4.0

1.2.34. Xrm Decrease Factor (XDM)

Definition: An ABR service parameter, XDF controls the decrease in ACR associated with Xrm. It is a power of two in range: [0, 1].

Discussion: For ABR traffic, XDM effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

1.2.35. Xrm

Definition: An ABR service parameter, Xrm limits the number of forward RM-cells which may be sent in the absence of received backward PM-cells. The range is 0-255.

Discussion: For ABR traffic, Xrm effects the time rate of change of the ACR. This value effects TCP round trip time calculations, which in turn effects TCP throughput.

Specification: AF-TM4.0

2. Performance Metrics

2.1. Definition Format (from RFC 1242)

Metric to be defined.

Definition: The specific definition for the metric.

Discussion: A brief discussion of the metric, its application and any restrictions on measurement procedures.

Measurement units: Intrinsic units used to quantify this metric. This includes subsidiary units; e.g., microseconds are acceptable if the intrinsic unit is seconds.

2.2. Definitions

2.2.1. ABR Rate Decrease Response Time (ARDRT)

Definition: The amount of time required by the SUT to adjust its transmission rate based on an ABR rate decrease request.

Discussion: During the ARDRT, cells transmitted by the SUT may be dropped by the network due to traffic policing. These dropped cells may contain a portion of an IP datagram. This may cause IP and TCP packet loss.

Measurement Units: seconds

2.2.2. ABR Rate Increase Response Time (ARIRT)

Definition: The amount of time required by the SUT to adjust its transmission rate based on an ABR rate increase request.

Discussion: During the ARIRT, the SUT will not fully utilize the available bandwidth. This will negatively impact IP and TCP throughput.

Measurement Units: seconds

2.2.3. RM-Cell Delay Variation (RM-CDV)

Definition: The variation in RM-cell transfer delay (RM-CTD) of RM-cells associated with a given traffic load, orientation and distribution, as well as an integration period. $RM-CDV = \max(RM-CTD) - \min(RM-CTD)$ where max and min indicate the maximum and minimum over the integration period, respectively.

Discussion: RM-CDV is a component of RM-cell transfer delay, induced by buffering and RM-cell scheduling.

RM-CDV effects the time required to notify the source of a change in the condition of the network. This in turn effects TCP round trip time calculations. Large values of RM-CDV will adversely effect TCP throughput and cause SAR timeout.

Measurement Units: seconds

2.2.4. RM-Cell Error Ratio (RM-CER)

Definition: The ratio of RM-cells with payload errors in a transmission in relation to the total number of RM-cells sent in a transmission associated with a given traffic load, orientation and

distribution, as well as an integration period. Note that errors occurring in the RM-cell header will cause RM-cell loss at the ATM layer. Note further that multiple errors in a payload will only be counted as one cell payload error.

$$\text{RM-CER} = \text{RM-Cells with payload errors} / \text{Total RM-Cells Transmitted.}$$

Discussion: The measurement is taken over a time interval and is desirable to be measured on an in-service circuit. RM-CER effects the time required to notify the source of a change in the condition of the network. This in turn effects TCP round trip time calculations. Large values of RM-CER will adversely effect TCP throughput and cause SAR timeout.

Measurement Units: dimensionless.

2.2.5. RM-Cell Loss Ratio (RM-CLR)

Definition: The ratio of lost RM-cells in a transmission in relation to the total RM-cells sent in a transmission associated with a given traffic load, orientation and distribution, as well as an integration period.

$$\text{RM-CLR} = \text{Lost RM-Cells} / \text{Total RM-Cells Transmitted.}$$

Discussion: The objective is to minimize RM-CLR. It is expressed as an order of magnitude, having a range of 10^{-1} to 10^{-15} and unspecified.

RM-CLR effects the time required to notify the source of a change in the condition of the network. This in turn effects TCP round trip time calculations. Large values of RM-CLR will adversely effect TCP throughput and cause SAR timeout.

Measurement Units: dimensionless.

2.2.6. RM-Cell Misinsertion Ratio (RM-CMR)

Definition: The ratio of RM-cells received at an endpoint that were not originally transmitted by the source end in relation to the total number of RM-cells properly transmitted associated with a given traffic load, orientation and distribution, as well as an integration period.

$$\text{RM-CMR} = \text{Misinserted RM-Cells} / \text{Total RM-Cells Transmitted.}$$

Discussion: The measurement is taken over a time interval and is desirable to be measured on an in-service circuit.

RM-CMR effects the time required to notify the source of a change in the condition of the network. This in turn effects TCP round trip time calculations. Large values of RM-CMR will adversely effect TCP throughput and cause SAR timeout.

Measurement Units: dimensionless.

2.2.7. RM-CRC Error Ratio

Definition: The ratio of RM-cells received at an endpoint which contain an invalid CRC in relation to the total number of RM-cells properly transmitted associated with a given traffic load, orientation and distribution, as well as an integration period.

Discussion: RM-CRC errors cause ATM RM-cells to be lost.

RM-CRC effects the time required to notify the source of a change in the condition of the network. This in turn effects TCP round trip time calculations. Large values of RM-CRC will adversely effect TCP throughput and cause SAR timeout.

Measurement Units: dimensionless

2.2.8. RM-Cell Transfer Delay (RM-CTD)

Definition: The elapsed time between a RM-cell exit event at the measurement point 1 (e.g., at the source UNI) and the corresponding RM-cell entry event at a measurement point 2 (e.g., the destination UNI) for a particular connection.

Discussion: The RM-cell transfer delay between two measurement points is the sum of the total inter-ATM node transmission delay and the total ATM node processing delay. This number is a constant and should not adversely effect performance.

Measurement units: seconds

2.2.9. Severely Errored Cell Block Ratio (SECBR)

Definition: The ratio of severely errored cell blocks in a transmission in relation to the total cell blocks sent in a transmission associated with a given traffic load, orientation and distribution, as well as an integration period.

SECBR = Severely Errored Cell Blocks/Total Transmitted Cell Blocks

Discussion: SECBR may cause the SUT to drop cells that may contain a portion of an IP datagram. This may cause IP and TCP packet loss.

Measurement Units: dimensionless.

3. Security Considerations

As this document is solely for providing terminology and describes neither a protocol nor an implementation, there are no security considerations associated with this document.

4. References

- [AF-TM4.0] ATM Forum, Traffic Management Specification Version 4.0, af-tm-0056.00, April 1996.
- [AF-TM4.1] ATM Forum, Traffic Management Specification Version 4.1, af-tm-0121.000, March 1999.
- [AF-UNI3.1] ATM Forum, User Network Interface Specification Version 3.1, September 1994.
- [AF-UNI4.0] ATM Forum, User Network Interface Specification Version 4.0, July 1996.

5. Editors' Addresses

Jeffrey Dunn
Advanced Network Consultants, Inc.
4214 Crest Place
Ellicott City, MD 21043 USA

Phone: +1 (410) 750-1700
EMail: Jeffrey.Dunn@worldnet.att.net

Cynthia Martin
Advanced Network Consultants, Inc.
4214 Crest Place
Ellicott City, MD 21043 USA

Phone: +1 (410) 750-1700
EMail: Cynthia.E.Martin@worldnet.att.net

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

