

Network Working Group
Request for Comments: 2761
Category: Informational

J. Dunn
C. Martin
ANC, Inc.
February 2000

Terminology for ATM 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 (2000). 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. The terms defined in this memo will be used in addition to terms defined in RFCs 1242, 2285, and 2544. This memo is a product of the Benchmarking Methodology Working Group (BMWG) of the Internet Engineering Task Force (IETF).

Introduction

This document provides terminology for benchmarking ATM based switching devices. It extends terminology already defined for benchmarking network interconnect devices in RFCs 1242, 2285, and 2544. 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 cell relay and signaling.

This memo contains two major sections: Background and Definitions. Within the definitions section is a formal definitions subsection, provided as a courtesy to the reader, and a measurement definitions sub-section, that contains performance metrics with inherent units. The divisions of the measurement sub-section follow the BISDN model.

The BISDN model comprises four layers and two planes. This document addresses the interactions between these layers and how they effect IP and TCP throughput. A schematic of the B-ISDN model follows:

	User Plane	Control Plane	
Services	IP	ILMI	UNI, PNNI
AAL	AAL1, AAL2, AAL3/4, AAL5	AAL5	SAAL
ATM	Cell Relay	OAM, RM	
Physical	Convergence		
	Media		

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). This document covers only CBR, VBR and UBR traffic types. ABR will be handled in a separate document. 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.

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. For the sake of clarity and continuity, this RFC adopts the template for definitions set out in Section 2 of RFC 1242. Definitions are indexed and grouped together in sections for ease of

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

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. These discussions pertain solely to the impact of a particular ATM parameter on IP or TCP; therefore, definitions which contain no configurable components or whose components will have the discussion: None.

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. ATM Adaptation Layer (AAL)

Definition: The layer in the B-ISDN reference model (see B-ISDN) which adapts higher layer PDUs into the ATM layer.

Discussion: There are four types of adaptation layers: AAL 1: used for circuit emulation, voice over ATM AAL2: used for sub-rated voice over ATM AAL3/4: used for data over noisy ATM lines AAL5: used for data over ATM, most widely used AAL type

These AAL types are not measurements, but it is possible to measure the time required for Segmentation and Reassembly (SAR).

Specification: I.363

1.2.2. ATM Adaptation Layer Type 5 (AAL5)

Definition: AAL5 adapts multi-cell higher layer PDUs into ATM with minimal error checking and no error detection. The AAL5 CPCS (Common Paer Convergence Sub-layer) PDU is defined as follows:

----- Higher Layer PDU -----	----- Padding (If needed) -----	----- Trailer -----
------------------------------------	---------------------------------------	---------------------------

Where the padding is used to ensure that the trailer occupies the final 8 octets of the last cell.

The trailer is defined as follows:

----- CPCS-UU -----	----- CPI -----	----- Length -----	----- CRC-32 -----
---------------------------	-----------------------	--------------------------	--------------------------

where:

CPCS-UU is the 1 octet Common Part Convergence Sub-layer User to User Indication and may be used to communicate between two AAL5 entities.

CPI is the 1 octet Common Part Indicator and must be set to 0.

Length is the 2 octet length of the higher layer PDU.

CRC-32 is a 32 bit (4 octet) cyclic redundancy check over the entire PDU.

Discussion: AAL5 is the adaptation layer for UNI signaling, ILMI, PNNI signaling, and for IP PDUs. It is the most widely used AAL type to date. AAL5 requires two distinct processes. The first is the encapsulation, on the transmit side, and de-encapsulation, on the receive side, of the higher layer PDU into the AAL5 CPCS PDU which requires the computation of the length and the CRC-32. The time required for this process depends on whether the CRC-32 computation is done on the interface (on-board) or in machine central memory (in core). On-board computation should produce only a small, constant delay; however, in core computation will produce variable delay, which will negatively effect TCP RTT computations. The second process is segmentation and re-assembly (SAR) which is defined below (see

SAR).

Specification: I.363.5

1.2.3. Asynchronous Transfer Mode (ATM)

Definition: A transfer mode in which the information is organized into 53 octet PDUs called cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.

Discussion: ATM is based on the ISDN model; however, unlike ISDN, ATM uses fixed length (53 octet) cells. Because of the fixed length of ATM PDUs, higher layer PDUs must be adapted into ATM using one of the four ATM adaptation layers (see AAL).

Specification: AF-UNI3.1

1.2.4. ATM Link

Definition: A virtual path link (VPL) or a virtual channel link (VCL).

Discussion: none.

Specification: AF-UNI3.1

1.2.5. ATM Peer-to-Peer Connection

Definition: A virtual channel connection (VCC) or a virtual path connection (VPC).

Discussion: none.

Specification: AF-UNI3.1

1.2.6. ATM Traffic Descriptor

Definition: A generic list of traffic parameters, which specify the intrinsic traffic characteristics of a requested ATM connection (see GCRA), which must include PCR and QoS and may include BT, SCR and best effort (UBR) indicator.

Discussion: The effects of each traffic parameter will be discussed individually.

Specification: AF-UNI3.1

1.2.7. ATM User-User Connection

Definition: An association established by the ATM Layer to support communication between two or more ATM service users (i.e., between two or more next higher entities or between two or more ATM-entities). The communications over an ATM Layer connection may be either bi-directional or unidirectional. The same Virtual Channel Identifier (VCI) is issued for both directions of a connection at an interface.

Discussion: Because ATM is connection oriented, certain features of IP (i.e. those which require multicast) are not available.

Specification: AF-UNI3.1

1.2.8. Broadband ISDN (B-ISDN) Model

Definition: A layered service model that specifies the mapping of higher layer protocols onto ATM and its underlying physical layer. The model is composed of four layers: Physical, ATM, AAL and Service.

Discussion: See discussion above.

Specification: I.321

1.2.9. Burst Tolerance (BT)

Definition: A traffic parameter, which, along with the Sustainable Cell Rate (SCR), specifies the maximum number of cells which will be accepted at the Peak Cell Rate (PCR) on an ATM connection.

Discussion: BT applies to ATM connections supporting VBR services and is the limit parameter of the GCRA. BT will effect TCP and IP PDU loss in that cells presented to an interface which violate the BT may be dropped, which will cause AAL5 PDU corruption. BT will also effect TCP RTT calculation. $BT = (MBS - 1) * (1/SCR - 1/PCR)$ (see MBS, PCR, SCR).

Specification: AF-TM4.0

1.2.10. Call

Definition: A call is an association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.

Discussion: none.

Specification: AF-UNI3.1

1.2.11. Cell

Definition: A unit of transmission in ATM. A fixed-size frame consisting of a 5-octet header and a 48-octet payload.

Discussion: none.

Specification: AF-UNI3.1

1.2.12. Call-based

Definition: A transport requiring call setups - see CALL definition.

Discussion: none.

Specification: AF-UNI3.1

1.2.13. Cell Delay Variation Tolerance (CDVT)

Definition: ATM layer functions may alter the traffic characteristics of ATM connections by introducing Cell Delay Variation. When cells from two or more ATM connections are multiplexed, cells of a given ATM connection may be delayed while cells of another ATM connection are being inserted at the output of the multiplexer. Similarly, some cells may be delayed while physical layer overhead or OAM cells are inserted. Consequently, some randomness may affect the inter-arrival time between consecutive cells of a connection as monitored at the UNI. The upper bound on the "clumping" measure is the CDVT.

Discussion: CDVT effects TCP round trip time calculations. Large values of CDVT will adversely effect TCP throughput and cause SAR timeout. See discussion under SAR.

Specification: AF-TM4.0

1.2.14. Cell Header

Definition: ATM Layer protocol control information.

Discussion: The ATM cell header is a 5-byte header that contains the following fields: Generic Flow Control (GFC) 4 bits Virtual Path Identifier (VPI) 8 bits Virtual Channel Identifier (VCI) 16 bits Payload Type (PT) 3 bits Cell Loss Priority (CLP) 1 bit Header Error Check (HEC) 8 bit CRC computed over the previous four octets

Each field is discussed in this document.

Specification: AF-UNI3.1

1.2.15. Cell Loss Priority (CLP)

Definition: This bit in the ATM cell header indicates two levels of priority for ATM cells. CLP=0 cells are higher priority than CLP=1 cells. CLP=1 cells may be discarded during periods of congestion to preserve the CLR of CLP=0 cells.

Discussion: The CLP bit is used to determine GCRA contract compliance. Specifically, two traffic contracts may apply to a single connection: CLP=0, meaning only cells with CLP=0, and CLP=0+1, meaning cells with CLP=0 or CLP=1.

Specification: AF-UNI3.1

1.2.16. Connection

Definition: An ATM connection consists of concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.

Discussion: none.

Specification: AF-UNI3.1

1.2.17. Connection Admission Control (CAC)

Definition: Connection Admission Control is defined as the set of actions taken by the network during the call set-up phase (or during call re-negotiation phase) in order to determine whether a connection request can be accepted or should be rejected (or whether a request for re-allocation can be accommodated).

Discussion: CAC is based on the ATM traffic descriptor (see ATM traffic descriptor) associated with the call as well as the presented and existing load. It may also be based on administrative policies such as calling party number required or access limitations. The effect on performance of these policies is beyond the scope of this document and will be handled in the BMWG document: Benchmarking Terminology for Firewall Performance.

Specification: AF-UNI3.1

1.2.18. Constant Bit Rate (CBR)

Definition: An ATM service category which supports a constant and guaranteed rate to transport services such as video or voice as well as circuit emulation which requires rigorous timing control and performance parameters. CBR requires the specification of PCR and QoS (see PCR and QoS).

Discussion: Because CBR provides minimal cell delay variation (see CDV), it should improve TCP throughput by stabilizing the RTT calculation. Further, as CBR generally provides a high priority service, meaning that cells with a CBR traffic contract usually take priority over other cells during congestion, TCP segment and IP packet loss should be minimized. The cost associated with using CBR is the loss of statistical multiplexing. Since CBR guarantees both throughput and CDV control, the connections must be subscribed at PCR. This is extremely wasteful as most protocols, e.g., TCP, only utilize full bandwidth on one half of a bi-directional connection.

Specification: AF-UNI3.1

1.2.19. Cyclic Redundancy Check (CRC)

Definition: A mathematical algorithm that computes a numerical value based on the bits in a block of data. This number is transmitted with the data, the receiver uses this information and the same algorithm to insure the accurate delivery of data by comparing the results of algorithm, and the number received. If a mismatch occurs, an error in transmission is presumed.

Discussion: CRC is not a measurement, but it is possible to measure the amount of time to perform a CRC on a string of bits. This measurement will not be addressed in this document. See discussion under AAL5.

Specification: AF-UNI3.1

1.2.20. End System (ES)

Definition: A system where an ATM connection is terminated or initiated. An originating end system initiates the ATM connection, and terminating end system terminates the ATM connection. OAM cells may be generated and received.

Discussion: An ES can be the user side of a UNI signaling interface.

Specification: AF-TEST-0022

1.2.21. Explicit Forward Congestion Indication (EFCI)

Definition: EFCI is an indication in the PTI field of the ATM cell header. A network element in an impending-congested state or a congested state may set EFCI so that this indication may be examined by the destination end-system. For example, the end-system may use this indication to implement a protocol that adaptively lowers the cell rate of the connection during congestion or impending congestion. A network element that is not in a congestion state or an impending congestion state will not modify the value of this indication. Impending congestion is the state when network equipment is operating around its engineered capacity level.

Discussion: EFCI may be used to prevent congestion by alerting a positive acknowledgement protocol and causing action to be taken. In the case of TCP, when EFCI cells are received the driver software could alert the TCP software of impending congestion. The TCP receiver would then acknowledge the current segment and set the window size to some very small number.

Specification: AF-TM4.0

1.2.22. Generic Cell Rate Algorithm (GCRA)

Definition: The GCRA is used to define conformance with respect to the traffic contract of the connection. For each cell arrival, the GCRA determines whether the cell conforms to the traffic contract. The UPC function may implement the GCRA, or one or more equivalent algorithms to enforce conformance. The GCRA is defined with two parameters: the Increment (I) and the Limit (L).

Discussion: The GCRA increment and limit parameters are mapped to CBR and VBR in the following fashion. For CBR, $I=1/PCR$ and $L=CDVT$ (CDV tolerance). For VBR, there are two GCRA algorithms running (dual leaky bucket). The first functions in the same fashion as CBR, $I=1/PCR$ and $L=CDVT$. The second, which polices cells which are in conformance with the first GCRA uses $I=1/SCR$ and $L=BT$ (see BT, CDV,

MBS, PCR and SCR).

Specification: AF-TM4.0

1.2.23. Generic Flow Control (GFC)

Definition: GFC is a field in the ATM header, which can be used to provide local functions (e.g., flow control). It has local significance only and the value encoded in the field is not carried end-to-end.

Discussion: none.

Specification: AF-UNI3.1

1.2.24. Guaranteed Frame Rate (GFR)

Definition: The GFR service provides the user with a Minimum Cell Rate (MCR) guarantee under the assumption of a given maximum frame size (MFS) and a given Maximum Burst Size (MBS). The MFS and MBS are both expressed in units of cells. GFR only applies to virtual channel connections (VCCs).

Discussion: GFR is intended for users who are either not able to specify the range of traffic parameters needed to request most ATM services, or are not equipped to comply with the (source) behavior rules required by existing ATM services. Specifically, GFR provides the user with the following minimum service guarantee: When the network is congested, all frames whose length is less than MFS and presented to the ATM interface in bursts less than MBS and at a rate less than PCR will be handled with minimum frame loss. When the network is not congested, the user can burst at higher rates.

The effect of GFR on performance is somewhat problematic as the policing algorithm associated with GFR depends on the network load; however, under congested condition and assuming a user who is following the GFR service agreement, it should improve performance.

Specification: AF-TM4.1

1.2.25. Header Error Control (HEC)

Definition: A check character calculated using an 8 bit CRC computed over the first 4 octets of the ATM cell header. This allows for single bit error correction or multiple bit error detection.

Discussion: none.

Specification: AF-UNI3.1

1.2.26. Integrated Local Management Interface

Definition: A management protocol which uses SNMPv1 carried on AAL5 to provide ATM network devices with status and configuration information concerning VPCs, VCCs, registered ATM addresses and the capabilities of ATM interfaces.

Discussion: ILMI is a conditionally required portion of UNI3.1; however, ILMI 4.0 has been issued as a separate specification. This document will refer to ILMI 4.0.

Specification: AF-ILMI4.0

1.2.27. Intermediate System (IS)

Definition: A system that provides forwarding functions or relaying functions or both for a specific ATM connection. OAM cells may be generated and received.

Discussion: An IS can be either the user or network side of a UNI signaling interface, or the network side of a PNNI signaling interface.

Specification: AF-TEST-0022

1.2.28. Leaky Bucket (LB)

Definition: Leaky Bucket is the term used as an analogous description of the algorithm used for conformance checking of cell flows from a user or network. See GCRA and UPC. The "leaking hole in the bucket" applies to the sustained rate at which cells can be accommodated, while the "bucket depth" applies to the tolerance to cell bursting over a given time period.

Discussion: There are two types of LB algorithms - single and dual. Single LB is used in CBR; dual LB is used in VBR (see CBR and VBR).

Specification: AF-TM4.0

1.2.29. Maximum Burst Size (MBS)

Definition: In the signaling message, the Burst Tolerance (BT) is conveyed through the MBS that is coded as a number of cells. The BT together with the SCR and the PCR determine the MBS that may be transmitted at the peak rate and still is in conformance with the GCRA.

Discussion: See the discussion under BT.

Specification: AF-TM4.0

1.2.30. Maximum Frame Size (MFS)

Definition: The MFS is the maximum length of a frame, expressed in units of cells, which in interface implementing GFR will accept during congested conditions (see GFR).

Discussion: During congestion, frames whose size is in excess of the MFS may be dropped or tagged. Assuming that the user is adhering to the MFS limit, this behavior should improve performance by improving congestion.

Specification: AF-TM4.1

1.2.31. Operations, Administration, and Maintenance (OAM)

Definition: A group of network management functions that provide network fault indication, performance information, and data and diagnosis functions.

Discussion: There are four types of ATM OAM flows: segment or end-to-end VP termination management (i.e. F4 segment, F4 E2E) and segment or end-to-end VC termination management (i.e. F5 segment, F5 E2E). These OAM cells can be used to identify fault management, connection verification, and loop back measurements.

Specification: AF-UNI3.1

1.2.32. Payload Type Indicator (PTI)

Definition: Payload Type Indicator is the Payload Type field value distinguishing the various management cells and user cells as well as conveying explicit forward congestion indication (see EFCI).

Example: Resource Management cell is indicated as PTI=110, End-to-end OAM F5 Flow cell is indicated as PTI=101.

Discussion: none.

Specification: AF-UNI3.1

1.2.33. Peak Cell Rate (PCR)

Definition: A traffic parameter, which specifies the upper bound on the rate at which ATM cells can be submitted to an ATM connection. This parameter is used by the GCRA.

Discussion: PCR directly limits the maximum data rate on an ATM connection. If a user violates the PCR, cells may be dropped resulting in Cell Loss. This in turn will negatively impact AAL5 PDUs, which may be carrying IP datagrams. See the discussion under SAR.

Specification: AF-TM4.0

1.2.34. Permanent Virtual Circuit (PVC)

Definition: This is a link with static route(s) defined in advance, usually by manual setup.

Discussion: none.

Specification: AF-UNI3.1

1.2.35. Permanent Virtual Channel Connection (PVCC)

Definition: A Virtual Channel Connection (VCC) is an ATM connection where switching is performed on the VPI/VCI fields of each cell. A permanent VCC is one that is provisioned through some network management function and left up indefinitely.

Discussion: none.

Specification: AF-UNI3.1

1.2.36. Permanent Virtual Path Connection: (PVPC)

Definition: A Virtual Path Connection (VPC) is an ATM connection where switching is performed on the VPI field only of each cell. A permanent VPC is one that is provisioned through some network management function and left up indefinitely.

Discussion: none.

Specification: AF-UNI3.1

1.2.37. Private Network-Network Interface (PNNI)

Definition: A routing information protocol that enables extremely, scalable, full function, dynamic multi-vendor ATM switches to be integrated in the same network.

Discussion: PNNI consists of signaling and routing between ATM network devices. PNNI signaling is based on UNI 4.0 signaling between two network side interfaces, while PNNI routing provides a mechanism to route ATM cells between two separate, autonomous ATM networks.

Specification: AF-PNNI1.0

1.2.38. Protocol Data Unit (PDU)

Definition: A PDU is a message of a given protocol comprising payload and protocol-specific control information, typically contained in a header. PDUs pass over the protocol interfaces that exist between the layers of protocols (per OSI model).

Discussion: In ATM networks, a PDU can refer to an ATM cell, multiple ATM cells, an AAL segment, an IP datagram and others.

Specification: Common Usage

1.2.39. Segmentation and Reassembly (SAR)

Definition: The process used by the AAL in the B-ISDN reference model (see B-ISDN) which fragments higher layer PDUs into ATM cells.

Discussion: SAR is not a measurement, but the speed in which SAR can be completed on a bit stream can be measured. Although this measurement is not included in this document, it should be noted that the manner in which SAR is performed will greatly effect performance. SAR can be performed either on the interface card (on board) or in machine central memory (in core). On-board computation should

produce only a small, constant delay; however, in core computation will produce variable delay, which will negatively effect TCP RTT computations. This situation is further complicated by the location of the CRC-32 calculation. Given an in core CRC-32 calculation, bus contention may cause on board SAR to be slower than in core SAR. Clearly, on board CRC-32 calculation and SAR will produce the most favorable performance results.

SAR performance will also be effected by ATM layer impairments. Cell error (CE), cell loss (CL), cell mis-insertion (CM) and cell delay variation (CDV) will all negatively effect SAR. CE will cause an AAL5 PDU to fail the CRC-32 check and be discarded, thus discarding the packet which the PDU contained. CL and CM will both cause an AAL5 PDU to fail the length check and be discarded. CL can have other effects depending on whether the cell which was lost is the final cell (PTI=1) of the AAL5 PDU. The following discussion enumerates the possibilities.

1. PTI=0 cell is lost. In this case, re-assembly registers a length discrepancy and discards the PDU.

2. PTI=1 cell is lost.

2. A. The AAL5 re-assembly timer expires before the first cell, PTI=0, of the next AAL5 PDU arrives. The AAL5 PDU with the missing PTI=1 cell is discarded due to re-assembly timeout and one packet is lost.

2. B. The first cell of the next AAL5 PDU arrives before the re-assembly timer expires. The AAL5 with the missing PTI=1 cell is prepended to the next AAL5 PDU in the SAR engine. This yields two possibilities:

2. B. i. The AAL5 re-assembly timer expires before the last cell, PTI=1, of the next AAL5 PDU arrives. The AAL5 PDU with the missing PTI=1 cell and the next AAL5 PDU are discarded due to re-assembly timeout and two packets are lost.

2. B. ii. The last cell of the next AAL5 PDU arrives before the re-assembly timer expires. In this case, AAL5 registers a length discrepancy and discards the PDU; therefore, the AAL5 PDU with the missing PTI=1 cell and the next AAL5 PDU are discarded due to their concatenation and two packets are lost.

2. C. Coupled with re-assembly, there exists some mechanism for identifying the start of a higher layer PDU, e.g., IP, and the cells associated with the first incomplete AAL5 PDU are discarded, resulting in the loss of one packet.

Specification: AF-UNI3.1

1.2.40. Sustainable Cell Rate (SCR)

Definition: The SCR is an upper bound on the conforming average rate of an ATM connection over time scales which are long relative to those for which the PCR is defined. Enforcement of this bound by the UPC could allow the network to allocate sufficient resources, but less than those based on the PCR, and still ensure that the performance objectives (e.g., for Cell Loss Ratio) can be achieved.

Discussion: SCR limits the average data rate on an ATM connection. If a user violates the SCR, cells may be dropped resulting in Cell Loss. This in turn will negatively impact AAL5 PDUs, which may be carrying IP datagrams. See the discussion under SAR.

Specification: AF-TM4.0

1.2.41. Switched Connection

Definition: A connection established via signaling.

Discussion: none.

Specification: AF-UNI3.1

1.2.42. Switched Virtual Channel Connection (SVCC)

Definition: A Switched VCC is one that is established and taken down dynamically through control signaling. A Virtual Channel Connection (VCC) is an ATM connection where switching is performed on the VPI/VCI fields of each cell.

Discussion: none.

Specification: AF-UNI3.1

1.2.43. Switched Virtual Circuit (SVC)

Definition: A connection established via signaling. The user defines the endpoints when the call is initiated.

Discussion: SVCs are established using either UNI signaling or PNNI signaling. The signaling state machine implements several timers, which can effect the time required for call establishment. This will effect TCP round trip time calculation, effecting TCP throughput. Specifically, there are two possibilities. In the case where Call Proceeding is not implemented, there is only one timer, T310, with a value of 10 seconds. In the case where Call Proceeding is implemented, there are two timers, T303 and T310, with the values 4 and 10 seconds, respectively. In either case, if a timer, either T303 or T310, expires after a Setup message is send, the calling party has the option of re-transmitting the Setup. In the T303 case, this yields a maximum setup time of 18 seconds and, In the T310 case, a maximum setup time of 20 seconds. Thus, the initial TCP RTT calculation will be on he order of 20 seconds.

Specification: AF-UNI3.1, AF-UNI4.0, AF-PNNI1.0

1.2.44. Switched Virtual Path Connection (SVPC)

Definition: A Switched Virtual Path Connection is one that is established and taken down dynamically through control signaling. A Virtual Path Connection (VPC) is an ATM connection where switching is performed on the VPI field only of each cell.

Discussion: none.

Specification: AF-UNI3.1

1.2.45. Traffic Contract

Definition: A specification of the negotiated traffic characteristics of an ATM connection.

Discussion: See discussions under BT, CAC, CDV, GCRA, PCR and SCR.

Specification: AF-TM4.0

1.2.46. Traffic Management (TM)

Definition: Traffic Management is the aspect of the traffic control and congestion control procedures for ATM. ATM layer traffic control refers to the set of actions taken by the network to avoid congestion conditions. ATM layer congestion control refers to the set of

actions taken by the network to minimize the intensity, spread and duration of congestion. The following functions form a framework for managing and controlling traffic and congestion in ATM networks and may be used in appropriate combinations.

- Connection Admission Control
- Feedback Control
- Usage Parameter Control
- Priority Control
- Traffic Shaping
- Network Resource Management
- Frame Discard
- ABR Flow Control

Discussion: See CAC and traffic shaping.

Specification: AF-TM4.0

1.2.47. Traffic Shaping (TS)

Definition: Traffic Shaping is a mechanism that alters the traffic characteristics of a stream of cells on a connection to achieve better network efficiency, while meeting the QoS objectives, or to ensure conformance at a subsequent interface. Traffic shaping must maintain cell sequence integrity on a connection. Shaping modifies traffic characteristics of a cell flow with the consequence of increasing the mean Cell Transfer Delay.

Discussion: TS should improve TCP throughput by reducing RTT variations. As a result, TCP RTT calculations should be more stable.

Specification: AF-UNI3.1

1.2.48. Transmission Convergence (TC)

Definition: A sub-layer of the physical layer of the B-ISDN model transforms the flow of cells into a steady flow of bits and bytes for transmission over the physical medium. On transmit the TC sublayer maps the cells to the frame format, generates the Header Error Check (HEC), and sends idle cells when the ATM layer has none. to send. On reception, the TC sublayer delineates individual cells in the received bit stream, and uses the HEC to detect and correct received errors.

Discussion: TC is not a measurement, but the speed in which TC can occur on a bit stream can be measured. This measurement will not be discussed in this document; however, its value should be constant and small with respect to cell inter-arrival at the maximum data rate.

Specification: AF-UNI3.1

1.2.49. Unspecified Bit Rate (UBR)

Definition: UBR is an ATM service category, which does not specify traffic related service guarantees. Specifically, UBR does not include the notion of a per-connection-negotiated bandwidth. No commitments are made with respect to the cell loss ratio experienced by a UBR connection, or as to the cell transfer delay experienced by cells on the connection.

Discussion: RFC 2331 specifies UBR service class for IP over ATM. UBR service models the "best effort" service type specified in RFC 791; however, UBR has specific drawbacks with respect to TCP service. Since UBR makes no guarantee with respect to cell loss (CL), cell delay variation (CDV) or cell mis-insertion(CM), TCP RTT estimates will be highly variable. Further, all negatively impact AAL5 re-assembly, which in turn may cause packet loss. See discussions under CDV and SAR.

Specification: AF-TM4.0

1.2.50. Usage Parameter Control (UPC)

Definition: Usage Parameter Control is defined as the set of actions taken by the network to monitor and control traffic, in terms of traffic offered and validity of the ATM connection, at the end-system access. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior, which can affect the QoS of established connections, by detecting violations of negotiated parameters and taking appropriate actions.

Discussion: See discussions under BT, CAC, CDV, GCRA, PCR and SCR.

Specification: AF-TM4.0

1.2.51. User-Network Interface (UNI)

Definition: An interface point between ATM end users and a private ATM switch, or between a private ATM switch and the public carrier ATM network; defined by physical and protocol specifications per ATM Forum UNI documents. The standard adopted by the ATM Forum to define connections between users or end stations and a local switch.

Discussion: none.

Specification: AF-UNI3.1

1.2.52. Variable Bit Rate (VBR)

Definition: An ATM Forum defined service category which supports variable bit rate data traffic with average and peak traffic parameters.

Discussion: VBR may potentially adversely effect TCP throughput due to large RTT variations. This in turn will cause the TCP RTT estimates to be unstable.

Specification: AF-TM4.0

1.2.53. Virtual Channel (VC)

Definition: A communications channel that provides for the sequential unidirectional transport of ATM cells.

Discussion: none.

Specification: AF-TM3.1

1.2.54. Virtual Channel Connection (VCC)

Definition: A concatenation of VCIs that extends between the points where the ATM service users access the ATM layer. The points at which the ATM cell payload is passed to, or received from, the users of the ATM Layer (i.e., a higher layer or ATM-entity) for processing signify the endpoints of a VCC. VCCs are unidirectional.

Discussion: none.

Specification: AF-TM3.1

1.2.55. Virtual Channel Identifier (VCI)

Definition: A unique numerical tag as defined by a 16 bit field in the ATM cell header that identifies a virtual channel, over which the cell is to travel.

Discussion: none.

Specification: AF-UNI3.1

1.2.56. Virtual Path (VP)

Definition: A unidirectional logical association or bundle of VCs.

Discussion: none.

Specification: AF-UNI3.1

1.2.57. Virtual Path Connection (VPC)

Definition: A concatenation of VPIs between Virtual Path Terminators (VPTs). VPCs are unidirectional

Discussion: none.

Specification: AF-TM3.1

1.2.58. Virtual Path Identifier (VPI)

Definition: An eight-bit field in the ATM cell header that indicates the virtual path over which the cell should be routed.

Discussion: none.

Specification: AF-UNI3.1

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. Physical Layer - SONET

2.2.1.1. Pointer Movements

Definition: Pointer Movements is the number of changes in a SONET pointer due to clock synchronization slips.

Discussion: SONET Pointer Movements can cause loss of information in the SONET payload envelop (SPE) which contains IP datagrams, either in the form of ATM cells or as PPP delimited PDUs.

Measurement Units: Per second.

2.2.1.2. Transport Overhead Error Count

Definition: SONET Transport Overhead Error Count is the number of SONET transport overhead errors detected.

Discussion: SONET Transport Overhead Errors cause SONET frames to be lost. These frames may contain IP datagrams; either in the form of cells or as PPP delimited PDUs.

Measurement Units: Positive integer

2.2.1.3. Path Overhead Error Count

Definition: SONET Path Overhead Error Count is the number of SONET path overhead errors detected.

Discussion: SONET Path Overhead Errors cause SONET frames to be lost. These frames may contain IP datagrams; either in the form of cells or as PPP delimited PDUs.

Measurement Units: Positive integer

2.2.2. ATM Layer

2.2.2.1. Cell Delay Variation (CDV)

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

Discussion: CDV is a component of cell transfer delay, induced by buffering and cell scheduling. Peak-to-peak CDV is a QoS delay parameter associated with CBR and VBR services. The peak-to-peak CDV is the $((1-a)$ quantile of the CTD) minus the fixed CTD that could be experienced by any delivered cell on a connection during the entire connection holding time. The parameter "a" is the probability of a cell arriving late.

CDV effects TCP round trip time calculations. Large values of CDV will adversely effect TCP throughput and cause SAR timeout. See discussion under SAR.

Measurement Units: seconds

2.2.2.2. Cell Error Ratio (CER)

Definition: The ratio of cells with payload errors in a transmission in relation to the total number of 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 cell header will cause cell loss at the ATM layer. Note further that multiple errors in a payload will only be counted as one cell payload error.

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

Discussion: The measurement is taken over a time interval and is desirable to be measured on an in-service circuit. CER is closely related to the number of corrupted AAL5 PDUs; however, there is not a direct numerical correlation between the number of errored cells and the number of corrupted AAL5 PDUs. There are two cases described below.

1. Only one cell in an AAL5 PDU contains payload errors. In this case, there is a one-to-one correspondence between cell payload errors and the number of corrupted AAL5 PDUs.

2. Multiple cells in the AAL5 PDU contain payload errors. In this case, there is not a one-to-one correspondence between cell payload errors and the number of corrupted AAL5 PDUs.

Measurement Units: dimensionless.

2.2.2.3. Cell Loss Ratio (CLR)

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

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

Discussion: CLR is a negotiated QoS parameter and acceptable values are network specific. The objective is to minimize CLR provided the end-system adapts the traffic to the changing ATM layer transfer characteristics. The CLR parameter is the value of CLR that the network agrees to offer as an objective over the lifetime of the connection. It is expressed as an order of magnitude, having a range of 10^{-1} to 10^{-15} and unspecified.

CLR indicates the number of ATM cells lost in relation to the total number of cells sent. CLR is closely related to the number of corrupted AAL5 PDUs; however, there is not a direct numerical correlation between the number of cells lost and the number of corrupted AAL5 PDUs. See the discussion under SAR.

Measurement Units: dimensionless.

2.2.2.4. Cell Misinsertion Ratio (CMR)

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

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

Discussion: The measurement is taken over a time interval and is desirable to be measured on an in-service circuit. CMR is closely related to the number of corrupted AAL5 PDUs; however, there is not a direct numerical correlation between the number of mis-inserted cells and the number of corrupted AAL5 PDUs. There are two cases described below.

1. Only one cell is mis-inserted into an AAL5 PDU. In this case, there is a one-to-one correspondence between cell mis-insertion and the number of corrupted AAL5 PDUs.

2. Multiple cells are mis-inserted into an AAL5. In this case, there is not a one-to-one correspondence between cell mis-insertion and the number of corrupted AAL5 PDUs.

Measurement Units: dimensionless.

2.2.2.5. Cell Rate Margin (CRM)

Definition: This is a measure of the difference between the effective bandwidth allocation and the allocation for sustainable rate in cells per second.

Discussion: This measures the amount of provisioned bandwidth which is not utilized. This lack of utilization may be caused by encapsulation overhead, e.g., AAL5 trailer and padding, or by the protocol itself, e.g., TCP usually transmits in only one direction.

Measurement units: Cells per second

2.2.2.6. CRC Error Ratio

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

Discussion: CRC errors cause ATM cells to be lost. Although this will appear as cell loss at the ATM layer, this measurement can be made in-service using a test probe which measures CRC errors at the TC layer.

Measurement Units: dimensionless

2.2.2.7. Cell Transfer Delay (CTD)

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

Discussion: The 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. While this number is a constant and should not adversely effect performance, it is a component in RTT.

Measurement units: seconds

2.2.3. ATM Adaptation Layer (AAL) Type 5 (AAL5)

2.2.3.1. AAL5 Re-assembly Errors

Definition: AAL5 Re-assembly Errors constitute any error, which causes the AAL5 PDU to be corrupted.

Discussion: AAL5 Re-assembly errors cause AAL5 PDUs to be lost. These PDUs may contain IP datagrams.

Measurement Units: Positive Integer

2.2.3.2. AAL5 Reassembly Time

Definition: AAL5 Reassembly Time constitutes the time between the arrival of the final cell in the AAL5 PDU and the AAL5 PDUs payload being made available to the service layer.

Discussion: AAL5 Reassembly time directly effects TCP round trip time calculations.

Measurement Units: seconds

2.2.3.3. AAL5 CRC Error Ratio

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

Discussion: AAL5 CRC errors cause AAL5 re-assembly errors. See discussion under AAL5 re-assembly errors.

Measurement Units: dimensionless

2.2.4. ATM Service: Signaling

2.2.4.1. CAC Denial Time

Definition: The amount of time required for CAC to determine that a call must be rejected.

Discussion: In the case where Call Proceeding is implemented, this number will be less than 4 seconds. Otherwise, it will be less than 10 seconds. Large values of this measurement will adversely effect performance on systems where an alternate, non-NBMA, service is available.

Measurement Units: seconds

2.2.4.2. Connection Establishment Time

Definition: The amount of time between the first Setup message from the calling party and the Connect message to the calling party.

Discussion: See discussion under SVC.

Measurement Units: seconds

2.2.4.3. Connection Teardown Time

Definition: The amount of between the Release message being sent and the Release Complete message being received.

Discussion: Large values of this measurement will adversely effect performance in systems where the total number of open calls or VCs is limited. Specifically, a new VC cannot be instantiated with the same VPI/VCI before the old one is released.

Measurement Units: seconds

2.2.4.4. Crankback Time

Definition: The amount of time between the issuance of the first release or release complete message by the switch where the current Designated Transit List (DTL) is blocked and the receipt of the SETUP with the updated DTLs by the target switch.

Discussion: This measurement does not take into account the amount of time associated with either the successful portion of the call setup transit or the time required for the calling party to receive .bp a response from the called party. As a result, the call may still fail to complete if the call setup timer on the calling party expires.

See discussion under SVC.

Measurement Units: seconds

2.2.4.5. Route Update Response Time

Definition: The amount of time between the receipt of a PNNI Topology State Element (PTSE), which is the PNNI routing PDU, containing a topology different from the current topology and the point at which the switch begins to generate DTLs reflecting the routing change.

Discussion: This measurement provides a lower bound on the amount of time during which SETUP messages will be forwarded along a sub-optimal or blocked path.

Measurement Units: seconds

2.2.5. ATM Service: ILMI

2.2.5.1. MIB Alignment Time

Definition: The amount of time between the issuance of the final cold start message and the final get response associated with the exchange of static MIB information.

Discussion: This measurement reflects the amount of time required by the switch and end system to exchange all information required to characterize and align the capabilities of both systems. It does not include address registration. It should also be noted that this measurement will depend on the number of MIB elements implemented by both systems.

Measurement Units: seconds

2.2.5.2. Address Registration Time

Definition: The amount of time between the initial set request issued by the switch and the final get response issued by the switch.

Discussion: This measurement assumes that the switch has checked the network prefix status object and the end system has checked the ATM address status object. In the case where the end system checks the ATM address status object only after the switch has issued a set request of the network prefix status object, this measurement will not reflect the actual time required to complete the address registration.

Measurement Units: seconds

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. Notices

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

5. References

- [AF-ILMI4.0] ATM Forum Integrated Local Management Interface Version 4.0, af-ilmi-0065.000, September 1996.
- [AF-TEST-0022] Introduction to ATM Forum Test Specifications, af-test-0022.00, December 1994.
- [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 (final ballot), btd-tm-01.02, July 1998.
- [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.

- [I.321] ITU-T, B-ISDN protocol reference model and its application, April 1991.
- [I.363] ITU-T, B-ISDN ATM Adaptation Layer Specification series, 1996-1997.
- [I.363.5] ITU-T, B-ISDN ATM Adaptation Layer Specification: Type 5 AAL, August 1996.

6. 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.
11241-B Skilift Court
Columbia, MD 21044 USA

Phone: +1 (410) 730-6300
EMail: Cynthia.E.Martin@worldnet.att.net

7. Full Copyright Statement

Copyright (C) The Internet Society (2000). 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.

