



Technology Backgrounder

ATM Environment

1.	Characteristics of ATM Service	1
2.	ATM Network Structure	1
3.	ATM Cell Structure	1
4.	ATM Network Operation	2
5.	Priorities in Traffic Handling.....	4
6.	ATM Network Service Categories.....	4
	ATM Network Transmission Parameters	5
	ATM Protocol Stack	6
7.	AAL Types Supported by RAD Products.....	8
8.	Operations and Maintenance	9
	OAM Flows	9
	Main Types of OAM Cells	10
9.	Frame Relay Support over ATM	11
	Frame Relay PDU Structure	11
	Handling of Frame Relay Traffic	13
10.	Circuit Emulation over ATM.....	15
	Protocol Stack.....	15
	Processing of E1 Data Stream	16
	Processing at the AAL1 SAR Sublayer	17
	Processing at the AAL1 Convergence Sublayer.....	17
	Dynamic Bandwidth Utilization (DBCES).....	19

1. Characteristics of ATM Service

ATM is a public high-performance cell-switched service. Its main characteristic is the use of very short packets, called “cells”, and intelligent switching nodes (“ATM switches”). The handling and routing of these cells, based on information carried in each cell header, is very simple and efficient, therefore high-speed transmission can be achieved.

Moreover, the use of short cells minimizes delay variation problems inherent in other packet transmission technologies, e.g., Frame Relay, etc., and simplifies cell handling.

ATM networks provide customers with individual access over various facilities, i.e., each individual subscriber has a dedicated physical access path, however the switching and transmission facilities are shared by the subscribers, in a way similar to the sharing of LAN and WAN resources.

2. ATM Network Structure

ATM service is usually provided by ASPs (ATM Service Providers) as well as data carrier service providers, and in general consists of several ATM switches interconnected by high-speed point-to-point transmission links.

Each transmission link interconnects two ATM switches. Time-division multiplexing on each transmission link is obtained by distributing the traffic generated by the subscribers over virtual paths (VPs). In accordance with the UNI specification, each link can have up to 256 VPs. Each VP is further subdivided into virtual channels (VCs). Each VP can support up to 65536 VCs.

The subscriber is connected by means of an access link to one of the network switches, and defines its communication requirements, e.g., average data rate, quality of service, etc., at subscription time. The connection point between the ATM network and the subscriber’s equipment is designated the User-Network Interface (UNI).

3. ATM Cell Structure

Figure 1 shows the structure of the ATM cell. Each cell has a 5-byte header that carries the control data, and a 48-byte information field, for a total of 53 bytes.

The header includes the following information fields:

- General flow control (GFC). The four GFC bits are always set to 0 when transmitted through the UNI, but can be used for local flow control purposes.
- Virtual path identifier (VPI). The VPI includes eight bits. When not all the bits are in use, the VPI is inserted in the least significant bit positions.
- Virtual circuit identifier (VCI). The VCI includes sixteen bits. When not all the bits are in use, the VCI is inserted in the least significant bit positions.

Note An ATM circuit is defined by a unique VPI.VCI pair (VCIs can be reused on a different VPs).

- Payload type (PT). The payload type, encoded by means of three bits, is used to discriminate between cells carrying user's payload, and the various types of cells used for ATM operations and maintenance (OAM). The OAM cells can be used for transfer of ATM signaling information, in-service testing (e.g., loopbacks) and performance monitoring, as well as to fill in idle intervals between cells carrying user's payload.
- Cell loss priority (CLP) indicator. This bit is used to indicate the priority set by the user for this cell: 0 for high cell priority, and 1 for low priority.
- Header error control (HEC). This field is designed for correction of single-bit, and detection of multiple-bit address errors.

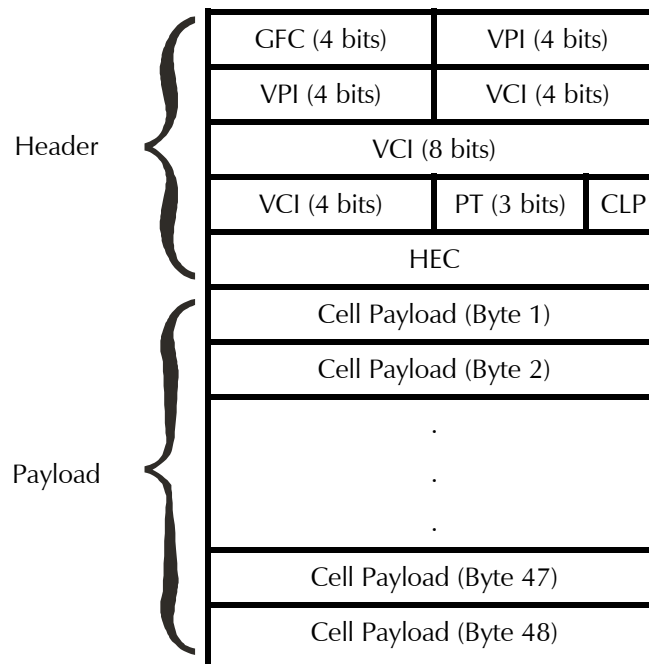


Figure 1. ATM Cell Structure

4. ATM Network Operation

Basically, ATM is a connection-oriented technology. ATM cells are transported over the network by setting up virtual channel connections (VCCs) between the UNIs of two subscribers wishing to communicate.

Note that the VCI in the ATM cell header is assigned per link, and therefore it may change across the network within the same VCC. Similarly, a VP groups VCs carried between two ATM entities, and therefore the VPI may also change along a given VCC.

Each ATM switch handles the VCs associated with a given VP as a block: such VCs are globally switched, without unbundling or processing the individual VCs in any way, and without changing their VCI numbers. This method ensures that the cell sequence of each VC is always preserved.

ATM network capacity design is based on traffic statistics, because the traffic handled by the ATM switches and the interconnecting links consists of large numbers of cells, generated by many independent sources. Statistically, the traffic generated by any specific set of users (e.g., users in a branch office) has a relatively constant rate, but its instantaneous rate varies rapidly.

FIFO buffers are used to temporarily store traffic peaks, and release data at a controlled rate. Therefore, the delay incurred by a cell in its passage through the network changes, because it depends on several factors:

- Link transmission delays. These delays decrease when the link data rates increase.
- Processing delays within the network and users equipment. These delays are relatively constant.
- Time spent by cells waiting in buffers. These delays vary with the traffic load.

Due to the bursty nature of the traffic, at times the instantaneous traffic offered by the network users may exceed the available ATM network resources (these resources include link bandwidth, buffer capacities, and the processing power of elements used to route the traffic).

Congestion occurs when the instantaneous traffic offered by the users exceeds the available network resources. When congestion occurs, some of the cells are lost (either because they are intentionally discarded within the ATM network, or because they are not accepted by the network).

Users must therefore agree with the network operators, at subscription time, on the amount of traffic that the network must transfer reliably. In principle, this traffic contract specifies the maximum traffic load that the network will accept from that particular user; as long as the user does not generate traffic in excess of the agreed-upon amount, the ATM network guarantees proper transfer.

Since the traffic generated by each user affects the service to other users, any ATM network, and each element connected to an ATM network must be able to perform traffic policing, to ensure that the traffic accepted from each user does not exceed the amount specified in its traffic contract (any excess traffic will be discarded).

Thus, the ATM network can be designed to offer to users a committed information rate, and may be able to support, when resources are available, higher peak rates.

5. Priorities in Traffic Handling

To enable a user to optimally utilize the ATM transmission capabilities subscribed for, it is necessary to allow each user to specify the relative importance (priority) of the traffic generated by its applications.

Two different types of priority can be implemented:

- Priority with respect to intentional discarding of cells in case of congestion. This priority is indicated by means of the CLP bit appearing in each cell header: lower-priority cells (cells with CLP = 1) are discarded preferentially in case of network congestion.

Note *The CLP bit in the cell header indicates the **loss priority**, therefore traffic which is more important to the user must be assigned low cell loss priority (CLP = 0).*

Note that if after discarding all the low priority cells, the offered traffic still exceeds the available network resources, the network will also discard priority cells (cells with CLP = 0), starting with the priority cells arriving from users that exceed their subscribed traffic rates.

- Priority with respect to the emission of cells toward the network. This priority is an internal priority mechanism whose purpose is to prevent the delaying of traffic with higher priority (as determined by user's preferences), when the equipment is already engaged in the transmission of lower-priority traffic. This is particularly important in equipment with multiple ports.

6. ATM Network Service Categories

As explained above, ATM network performance depends on traffic load versus network resources, therefore it is not constant. Users nevertheless need a predictable quality of service (QoS).

The required QoS depends on the application using the ATM connection, and therefore it must be clearly defined (this is done at subscription time, by means of the traffic contract, which is the collection of traffic parameters agreed upon between the user and the ATM network operator). For example:

- An ATM connection used for video or voice transmission generates traffic at a constant rate, and requires that the transmission delay through the network be also relatively constant (to a certain degree, buffers can be used to smooth out variations). On the other hand, the occasional loss of a few cells could be tolerable, because of the subjective interpretation of video and audio information.
- An ATM connection used for file transmission, LAN interconnection, or data communications, is not very sensitive to delay: such connections can tolerate any delay up to the delay that causes time-out of higher-level protocols, but is

very sensitive to the loss of cells (any cell lost, or received with errors, requires retransmission of the whole message).

For standardization, the ATM Forum standards specify several types of service categories:

- CBR** Constant bit rate service, intended for real-time applications that require a constant bandwidth, and tightly constrained delay and delay variation. The required bandwidth must be available as long as the connection is active. A typical application for CBR service is circuit emulation.
- VBR** Variable bit rate service, intended for applications that require constraints on the maximum delay and delay variation, e.g., video, voice, etc. Two types of VBR service are defined:
- RT-VBR – real-time VBR, which requires tight control on delay and delay variation.
 - NRT-VBR – non-real-time VBR, with more relaxed requirements).
- UBR** Unspecified bit rate service, intended for non-real-time applications which do not require tightly constrained delay and delay variation, e.g., data communications applications such as file transfer, e-mail, etc.
- UBR+** Same as UBR, except that it enables the user to specify a minimum desired cell rate.
- ABR** Available bit rate service. This service is similar to UBR, except that it is capable of adaptation to changes in ATM transfer capabilities without inducing excessive cell loss.

ATM Network Transmission Parameters

To enable consistent definition of transmission requirements, a standard set of traffic parameters has been defined in the ATM Forum traffic handling specifications. The traffic parameters include:

- Peak Cell Rate (PCR)** Maximum rate of cells accepted from the user. Generally, cells received at rates exceeding the PCR are discarded. PCR values generally specify the maximum rate for all the cells (CLP=0+1), however sometimes separate PCR values may be specified for cells with CLP=0 (high priority) and CLP=0+1 (all cells).
- Maximum Cell Transfer Delay (maxCTD)** Maximum time allowed for the transfer of a cell to its destination. Cells spending a longer time in transit are considered useless to the user's application.

Peak-to-Peak Cell Delay Variation (CDV)	The maximum tolerable variation in the cell transit time through the network.
Sustainable Cell Rate (SCR)	The long-term average cell rate provided by the user.
Maximum Burst Size (MBS)	The maximum number of cells allowed in a single burst. Longer bursts may be discarded.
Cell Loss Ratio (CLR)	The maximum ratio of cells that may get lost in the network when the offered user's traffic conforms to the traffic contract.
Minimum Desired Cell Rate (MDCR)	The minimum desired rate of cells for the UBR+ mode.

A traffic contract includes the parameters listed above (or at least the mandatory parameters, e.g., PCR). The UNI function is to enforce the traffic contract on the traffic arriving from the user, to ensure that it does not exceed the agreed-upon limits (any non-conforming traffic is selectively discarded in accordance with the principles listed above, to ensure that users exceeding their limits do not degrade the service to conforming users).

ATM Protocol Stack

The ATM transmission method uses very short cells, whereas most applications use much longer service data units (SDUs): for example, Ethernet frames have lengths up to approximately 1500 bytes, Frame Relay frames can be more than 9000 bytes long, etc.

Therefore, it is necessary to define a set of procedures (stack of protocols) that will enable inserting and removing application data from ATM cells, and convert the address information to enable routing to the desired destination.

Figure 2 shows the general structure of the ATM protocol stack.

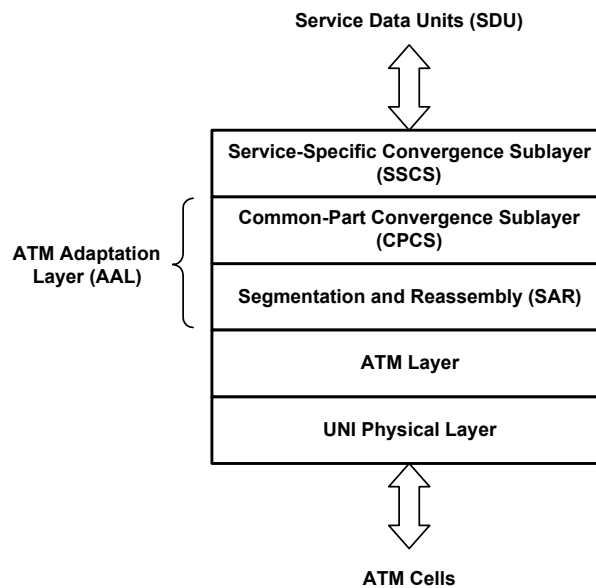


Figure 2. ATM Protocol Stack

The general structure of the ATM protocol stack is as follows:

- The top layer is the service-specific convergence sublayer (SSCS), which converts the service data units received from the users application to the CPCS format. The SSCS depends on the users application type.
- The insertion of payload data into the 48-byte information field of an ATM cell is accomplished by the third layer, the ATM Adaptation Layer (AAL). The function of the AAL is to map the protocol data units (PDUs) of higher layers into the ATM layer. AAL functions are specified in ITU-T Rec. I.363.

Using an AAL gives ATM the flexibility to carry entirely different types of service within the same transmission format, because the AAL is not an ATM network process, but instead is performed by the network terminating equipment.

Thus, after data is processed by an AAL type that can appropriately handle the users application characteristics, the only task an ATM network must perform is to route cells from point to point, in accordance with the information in the cell headers.

The AAL layer consists of two sublayers:

- The top sublayer is the Common Part Convergence Sublayer (CPCS), which provides the service access point (SAP) to the ATM service.
- The second sublayer is the segmentation and reassembly (SAR) sublayer, which splits the CPCS PDUs into 48-byte units, that are then passed to the ATM layer for insertion in ATM cells.
- The next layer, located atop the UNI physical layer, is the ATM layer.

- The lowest layer in the ATM protocol stack is the UNI physical layer. This is the layer responsible for the transfer of ATM cells by the uplink port. For example, for an uplink port with E1 interface, the functions of the physical layer are:
 - Generation of the E1 frame and multiframe structure, and the corresponding F-bit patterns.
 - Insertion of ATM cells in the E1 frame timeslots. Only 30 timeslots are used (timeslots 0 and 16 are not used for ATM cell transmission). ATM cells are byte-aligned to the timeslot boundaries (each payload byte of an ATM cell is inserted in one timeslot).
 - Transmission and reception of E1 signal.

7. AAL Types Supported by RAD Products

Several AAL types have been standardized, to provide optimal support for various types of payload data.

Some of the RAD's products support the following AAL types:

- AAL1, which is intended for constant bit rate (CBR) applications. In CBR applications, e.g., circuit emulation service (CES) for E1 signal transport, etc., it is necessary to emulate the transmission characteristics of regular data transmission circuits, and in particular fixed delay and preservation of timing accuracy.

For example, the data transferred through an E1 user port is handled by AAL1. Analog voice can also be transmitted over AAL1.

- AAL2, which is intended for variable bit rate (VBR) applications that require maintaining the timing relationship between the source and the destination, but the bit rate may vary. For example, digitized voice can be supported by AAL2, using the Loop Emulation Service (LES). In particular, AAL2 must be used to support Emulated Loop Control Protocol (ELCP).
- AAL5, which is intended for connectionless and connection-oriented variable bit rate (VBR) services. AAL5 is primarily designed to meet the communication requirements of LANs and other similar applications. The synchronous data is handled by AAL5.

8. Operations and Maintenance

The ATM standards specify the methods to be used for operations and maintenance (OAM) in ATM networks. The OAM functions to be made available in ATM networks include:

- Performance monitoring – enables monitoring the normal operation of equipment and links, and detection of defects and failures.
- Detection of defects and failures – made by identification of abnormal performance.
- System protection – enables correction of defect and failure effects by selection of alternate facilities.
- Collection of failure and performance information – enables management stations to obtain information on system state.
- Support for fault localization – enables use of internal or external test functions to detect faults whose location is not obvious from the failure and performance information normally generated by the system.

OAM Flows

The applicable standards, e.g., ITU-T Rec. I.610, define five **OAM flows**. An OAM flow is a bidirectional exchange of OAM information (generally performed by means of a special type of cells, called OAM cells) between peer entities. The five OAM flows are designated F1 through F5, where the flows F1, F2, and F3 cover the physical layer, and the flows F4 and F5 cover the ATM layer.

The physical layer OAM flows, used to detect physical layer failures and enable the management system to monitor transmission performance, are as follows:

- F1** Lowest-level OAM flow, covers a portion (called regenerator section) of an ATM system located between two entities directly connected by a physical link.
- F2** Covers a digital section of an ATM system. A digital section is located between two end points, interconnected by one or more physical links.
- F3** Covers a transmission path in an ATM system. The transmission path is delineated by end points that can assemble/disassemble cells, and process the information carried in the cells, e.g., can detect cell boundaries, analyze headers, and perform OAM functions. A transmission path consists of one or more digital sections.

The ATM layer OAM flows, used to monitor the operation and performance of the ATM layer, are as follows:

- F4** VP-level OAM flow, covers a VP connection. A VP connection is terminated at end points that can terminate virtual path connections. The virtual path connection consists of one or more transmission paths.
- F5** Highest-level OAM flow, at the VC connection level. A VC connection

is terminated at end points that can terminate virtual channel connections. The virtual channel consists of one or more virtual path connections.

The F5 flow is used for end-to-end fault management, for maintenance, and for performance management over a virtual channel connection.

The F1, F2, F3, and F5 OAM flows in accordance with ITU-T Rec. I.610.

Main Types of OAM Cells

The OAM cells are special cells that are transmitted in idle intervals, during which no payload waits to be transmitted (usually, forced insertion is used if no idle interval is available within a group of 65535 cells). The cell type and function is identified by one byte, sent immediately after the 5-byte ATM cell header; the remaining bytes may carry additional information to support the cell function.

Fault management uses the following types of OAM cells:

- CC Cells** CC (continuity check) cells are periodically transmitted downstream (nominally once per second) by a VC end point when no user cells are received for three consecutive seconds, and stops after one valid user cell is received.
- AIS Cells** AIS (alarm indication signal) cells are periodically transmitted downstream (nominally once per second) by a VC end point which detects a VC failure (unavailability). The transmission of AIS cells at the VC level starts when neither user cells, nor continuity check cells are received for three consecutive seconds, and stops after one valid user or continuity check cell is received.
- The AIS cell may include information on the type and location of the failure.
- FERF Cells** FERF (far-end report failure) cells are periodically transmitted upstream (nominally once per second) by a VC end point which started transmitting AIS cells, or detected a VC failure (unavailability).
- The FERF cell may include information on the type and location of the failure.

Note For Frame Relay operation, the corresponding Frame Relay OAM functions (link integrity verification, and the AIS and RDI (remote defect indication) conditions) are mapped from the Frame Relay side to the ATM side in accordance with the applicable Frame Relay Forum documents:

- Document FRF.5 for ports configured for the network interworking mode.
 - Document FRF.8 for ports configured for the service interworking mode.
-

The payload field of performance monitoring cells carry data which enables error detection. The main fields are as follows:

- PM cell sequence number (0 to 256), followed by the total number of user cells sent since the last performance monitoring cell has been transmitted (0 to 65535 cells).
- Block error detection code, carries a checksum that enables detection of errors in the payload fields of the users cells preceding the performance monitoring cell.

Maintenance cells carry commands:

- PM activate/deactivate command: starts/stops the transmission and evaluation of performance monitoring cells.
- Continuity check activate/deactivate command. The command instructs the receiving equipment whether to transmit and/or receive (evaluate) CC cells.
- Loopback activate/deactivate command. The loopback controlled by this command is performed at the ATM layer on OAM cells, and is used to instruct its destination (one or both of the VCC end points, or any other equipment that can evaluate OAM cells) to return OAM cells toward the command source (the equipment that sent the loopback command).

9. Frame Relay Support over ATM

To transfer Frame Relay traffic over the ATM network, communication devices can be configured either to encapsulate the Frame Relay PDUs (frames) in the ATM protocol, or to translate them to the ATM protocol.

Frame Relay PDU Structure

Figure 3 shows the structure of the Frame Relay frames (formally, the frames are called protocol data units – PDUs).

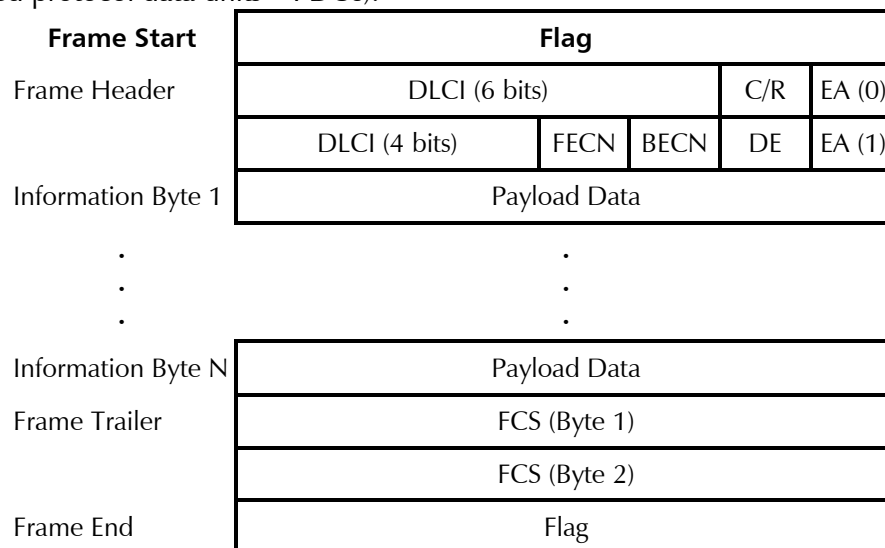


Figure 3. Frame Relay PDU (Frame) Structure

Each Frame Relay PDU starts and ends with an one-byte flag, and in addition includes:

- Two-byte header that carries the routing control data.
- Information field, which carries payload data of variable length.
- Two-byte trailer.

The frame fields have the following functions:

- **Flag.** The flag is used to separate between consecutive frames. Only one flag is needed to separate consecutive frames, therefore the ending flag of a frame can also serve as the opening flag of the next frame.
- **Data Link Connection Identifier (DLCI).** The standard DLCI consists of 10 bits, therefore the allowed range of DLCIs is 0 through 1023. Some values in this range are reserved. The reserved DLCIs are as follows:

DLCI	Function
0	Used for local in-channel signaling
1 through 15	Reserved for network use
16 to 991	Available for users traffic routing
992 to 1023	Reserved for network use

The function of the DLCI is to serve as a routing label for frames.

- **Command/Response (C/R) bit.** The C/R bit is transparently transferred, and is available for users protocols.
- **Extended Address (EA) bit.** The EA bit is used to indicate the use of extended addressing. In general, only 10-bit DLCIs are used, and therefore the EA bit is set to 0 in the first header byte, and to 1 in the second header byte.
- **Forward Explicit Congestion Notification (FECN) bit.** The FECN bit is normally set to 0. When congestion occurs, that is, when the network does not have sufficient resources to handle the users traffic at the current rate, the FECN bit can be set by the network to 1, to notify downstream equipment (that is, other equipment that receive the frame) that congestion avoidance procedures should be performed.

This bit is intended for use by protocols which enable the destination to control the traffic generated by the source: in response to the reception of frames with the FECN indicator set, the destination is expected to signal the source that it must reduce the rate at which it sends data.

- **Backward Explicit Congestion Notification (BECN) bit.** The BECN bit is normally set to 0. When congestion occurs, the BECN bit can be set by the network to 1, to notify upstream equipment (that is, the equipment that transmits the frames) that congestion avoidance procedures should be performed.

This bit is intended for use by protocols where the traffic is controlled by the source: in response to the reception of frames with the BECN indicator set, the source is expected to reduce the rate at which it sends data.

- **Delete Eligibility (DE) indication.** This bit is used to indicate the priority set by the user for this frame. This bit is set to 1 for frames that have lower priority and may be discarded in case of congestion, or when the network rejects users data because the user exceeds its subscribed data rate.
- **Information field.** The information field carries the users data.
- **Frame check sequence (FCS).** The FCS field consists of two bytes, calculated in accordance with the CRC16 polynomial specified by ITU-T Rec. Q.921. The CRC32 polynomial can be also supported.

Note *The CRC16 polynomial is effective at detecting errors in frames up to 4096 bytes long. For longer frames, it is necessary to select the CRC32 polynomial. If some equipment on the transmission path does not support the CRC32 polynomial, the user must select higher-level protocols that can provide the error detection function.*

Handling of Frame Relay Traffic

When using the Frame Relay interface mode, the ATM protocol stack includes an additional sublayer, located atop the AAL5 CPCS, as shown in [Figure 4](#).

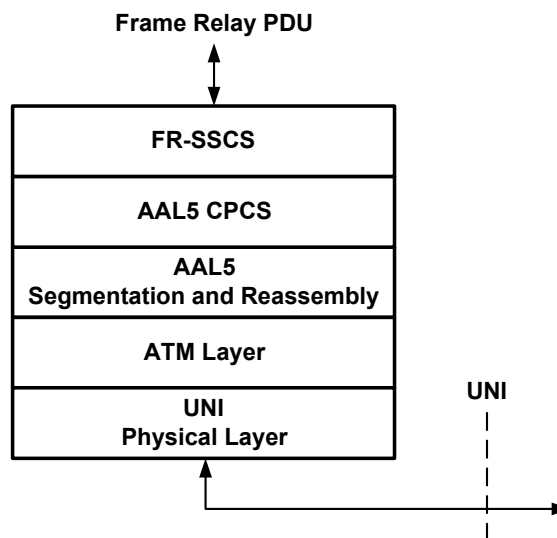


Figure 4. ATM Protocol Stack with FR-SSCS

This sublayer, called Frame Relay Service Specific Convergence Sublayer (FR-SSCS), performs the functions required to transport the Frame Relay PDU through the ATM network. In particular, it is necessary to perform the following functions:

- Convert the Frame Relay DLCI to the desired ATM VPI.VCI pair. For each DLCI the user can specify the desired VPI.VCI pair.

In general, each DLCI is routed through a unique VPI.VCI pair (ATM connection), however for connections defined on Frame Relay ports configured for the network interworking mode it is possible to route several DLCIs through the same ATM connection.

- Convert between the bits used to control congestion, in accordance with the documents applicable to the Frame Relay port interworking mode (Frame Relay Forum document FRF.5 for ports configured for the network interworking mode, and document FRF.8 for ports configured for the service interworking mode):
 - Conversion between the congestion indications supported by the ATM network to the FECN and BECN bits of the Frame Relay protocol.
 - Conversion between the DE bit received from the Frame Relay DTE and the CLP bit in the corresponding ATM cells. In this respect, the user can perform conversion of equivalent indications, or to set the bits to a predetermined value.
- Convert OAM information between the Frame Relay side and the ATM side in accordance with the documents applicable to the Frame Relay port interworking mode (FRF.5 and FRF.8).

The user can select additional conversion services to be performed, in accordance with the selected interworking method:

- For network interworking, the user can specify the DLCI to be used at the far end of the link, to complete the end-to-end connection.
- For service interworking, the user can select the handling of the upper layer user protocols:
 - Multiprotocol translation - converting the encapsulation header from the format required by RFC 1490 (Frame Relay) to the format of RFC 1483 (ATM).

This mode is used for transferring internetworking (routed or bridged) protocols, e.g., in LAN-to-LAN applications.

- Transparent - encapsulation header is transferred transparently.

This mode is used when the users protocols do not conform with the protocols supported in the translation mode, but are compatible with the users equipment at the two ends of the connection. A typical application for this mode is packetized voice.

10. Circuit Emulation over ATM

The circuit emulation service (CES) enables the transmission of constant bit rate (CBR) data over the ATM network, in accordance with the ATM Forum Circuit Emulation Service Interoperability Specification. CES uses AAL1.

Circuit emulation service is usually supported using AAL1 in accordance with ITU-T Rec. I.363. In addition, analog voice traffic can also be supported when AAL1 is used.

Figure 5 illustrates the structure of the ATM protocol stack used for CBR transmission using AAL1.

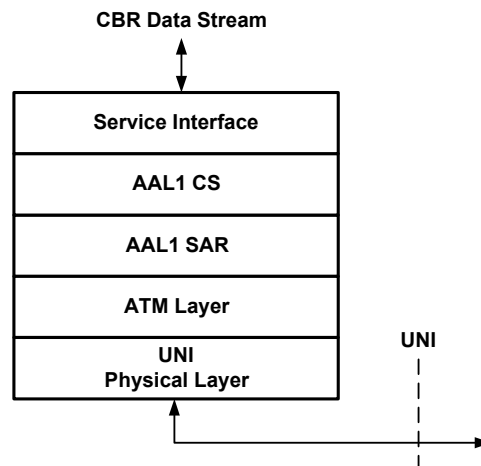


Figure 5. ATM Protocol Stack for CBR Transmission

Protocol Stack

A general description of the ATM protocol stack structure is given in the

Characteristics of ATM Service section above. The protocol stack for the CBR service includes the following layers:

- Service Interface. The service interface is the user port, and its main function is the processing of the E1 data stream received from users equipment.
- The next layer is the AAL1 convergence sublayer, followed by the AAL1 SAR sublayer.
- The lower layers (ATM layer and UNI physical layer).

The following sections explain the processes, which are specific to the AAL1 protocol stack.

Processing of E1 Data Stream

The processing of the incoming E1 data stream includes:

- Physical layer processing:
 - Regeneration of incoming signal and recovery of the associated timing.
 - Transmission of E1 signal to the line.
 - Detection of physical layer defects and failures, support for test loopbacks, and collection of physical layer performance statistics.
- Port timing. The service interface also handles the port timing.

In particular, the service interface generates the transmit-to-user timing. This timing is locked to the timing recovered from the receive signal at the remote end of the CBR link, otherwise bit integrity will be lost. This function is also supported by the CS sublayer.

- Synchronization to the frame and super-frame structure, in accordance with the selected E1 framing mode.
- Retrieval of payload data from the specified timeslots. For each connection defined on the E1 port, the user can select the timeslots whose data is to be transmitted to the remote end through the ATM connection.

When the connection to the ATM network is also made through an E1 link, it is not possible to transmit the whole payload, because some of the main link bandwidth is reserved for ATM overhead: the maximum number of timeslots that can be transmitted is 25 for E1 uplinks, and 18 f uplinks.

- Signaling information processing. The user can select between two modes:
 - Basic (non-CAS) service: in this mode, no special processing of signaling information is performed. Signaling information may be transmitted in a separate timeslot (for example, using CCS - common channel signaling).
 - CAS mode: in this mode, the service interface processes the timeslot 16 or "robbed-bit" signaling information associated with the connection timeslots, and transfers it using AAL1 services.

The user payload data, and optionally the associated signaling data, is transferred to the AAL1 convergence sublayer.

Processing at the AAL1 SAR Sublayer

In the transmit direction, the AAL1 SAR sublayer receives data from the AAL1 convergence sublayer in CS PDUs (CS protocol data units) containing 47 bytes, and generates 48-byte SAR PDUs which are transferred to the ATM layer. The reverse is performed in the receive direction.

The user bytes are aligned with the ATM cell bytes.

The SAR PDU is formed by adding a SAR header of one byte to the 47 bytes of a CS PDU. The SAR header contains the following information:

- Sequence number (SN) field (four bits). The sequence number is in the range of 0 to 7, and it enables keeping track of the cell order at the receiving end. This enables the detection of lost, or misinserted cells.

In addition, one bit indicates whether a CS sublayer is present (for structured data transfer, the use of the CS sublayer is mandatory).

- Sequence number protection (SNP) field (four bits). The SNP field includes three bits that carry a checksum generated by a third-order CRC code (CRC-3), followed by an even parity bit. The parity bit improves the error detection capabilities. Since some equipment does not support the parity bit, the user can disable its use.

The SNP field enables the correction of single errors in the SAR header, and the detection of two errors.

Processing at the AAL1 Convergence Sublayer

The AAL1 functionality is defined in ITU-T Rec. I.363.1. The function of the CS sublayer is to enable the transmission of fixed-size data blocks (where each block consists of an integer number of bytes), in the payload fields of ATM cells, while detecting missing or misinserted cells, and providing source clock recovery.

The start of each block is indicated by a special type of pointer, the AAL1 structure pointer. The AAL1 structure pointer consists of one byte and is inserted in the first byte of the payload field in each even cell (0, 2, 4, 6) in the eight-cell sequence defined by means of the SN field. It indicates the offset, in bytes, between the end of the pointer field, and the beginning of the next structured block in the 93-byte payload part.

The AAL1 convergence sublayer inserts the data from the service interface in CS-PDUs (CS protocol data units) containing 47 bytes. Therefore, the payload field of cells carrying the AAL1 structure pointer may include a maximum of 46 payload bytes.

CS PDU structure depends on the service type: basic (non-CAS) or CAS service.

Basic Service CS PDU

For the basic service, the CS PDU includes only payload data. The data block is formed by collecting N bytes (one from each timeslot to be carried by the connection) from each frame. N is the data rate multiplier, $N \times 64$ kbps, in the range of 1 through 26 for E1 uplinks and 1 through 18 uplinks.

Figure 6 shows a typical data block for $N = 3$.




Figure 6. Basic Service - Data Block Structure

Note that the time required to collect enough timeslots to fill a cell depends on the number of timeslots being carried. At the minimum data rate supported by the user port (128 kbps), corresponding to two timeslots per block, the time required to fill a cell is approx. 3 msec. To enable the user to reduce this inherent delay, it is possible specifying the number of payload bytes per cell: the remaining cell bytes are then filled with padding bytes.

CAS CS PDU

For the CAS service, the CS PDU includes a payload data part and a signaling substructure, used to carry the signaling information associated with the connection timeslots.

- The payload part of the data block has the length of one multiframe (that is, 16 E1 frames).
- The signaling substructure contains the signaling bits associated with the payload timeslots for the current multiframe. With 256S framing, the signaling substructure contains four signaling bits (A, B, C, D) for each payload timeslot.

Padding bits (0s) are used to fill fields which do not carry signaling payload.

Figure 7 shows a typical data block for $N = 3$ when 256S framing is used.

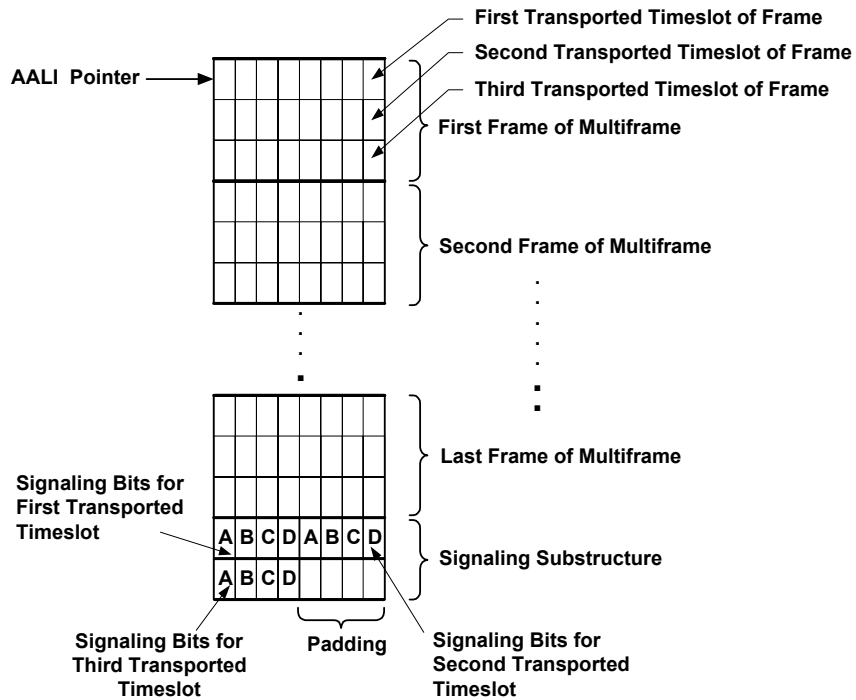


Figure 7. CAS Service - Data Block Structure

Dynamic Bandwidth Utilization (DBCES)

The principles governing the building of the data block structures described above for the CAS service can be extended to support additional signaling capabilities. One of the most important applications is dynamic bandwidth utilization over CES (DB-CES).

To enable dynamic bandwidth utilization in an ATM network, it is necessary to detect which timeslots of a given TDM trunk are active and which are inactive. When an inactive state is detected in a specific timeslot, the timeslot is dropped from the next ATM structure and the bandwidth it was using may be reutilized for other services. This method may be applied utilizing any method of timeslot activity detection, e.g., CAS, CCS.

This approach preserves the CES transmission characteristics while enabling efficient utilization of the available ATM link bandwidth.