SG24-4958-00

IBM Networked Video Solution Over ATM Implementation

May 1997

International Technical Support Organization
Raleigh Center

International Technical Support Organization

SG24-4958-00

IBM Networked Video Solution Over ATM Implementation

May 1997

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, "Special Notices" on page 183.

First Edition (May 1997)

This edition applies to the 8300 Video Access Node (VaN) and 8260 Video Distribution Module (VDM). Overview information is also provided on the IBM MediaSTREAMER.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

Contents

Preface

Intense progress in digital signal processing and transmission media in recent years has increased the popularity and use of multimedia systems.

From analog video to digital audiovisual services, this redbook describes the concepts of networked video and its implementations. A detailed explanation of audio and video compression techniques is presented, as well as the principles of Asynchronous Transfer Mode (ATM) and its use to transport compressed video, which will help the reader to understand this new area in the multimedia environment.

This redbook presents IBM solution examples for high-quality networked video services and provides detailed instructions for implementations using the 8300 Video Access Node Version 1, the 8260 Video Distribution Module Version 1 and the MediaSTREAMER Version 1, which will help the reader to quickly understand, market and support the IBM video over ATM solutions.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

Ricardo Haragutchi is a Senior ITSO Specialist for Networking, Internet and Multimedia at the Systems Management and Networking ITSO Center, Raleigh. He holds a Bachelors of Science degree in Electrical Engineering from Escola Politecnica in Sao Paulo University. He writes extensively and teaches IBM classes worldwide on areas such as routing, remote access, and Internet environment. Before joining the ITSO 2 years ago, Ricardo worked in the Field Systems Center (FSC) in IBM Brazil as Senior System Engineer.

Alvaro Brito is a systems and product specialist in Brazil. He has six years of experience in the networking field. He holds a degree in Electronic Engineering from the Pontificia Universidade Catolica do Rio de Janeiro - PUC/RJ. His areas of expertise include the design of network solutions using technologies such as time division multiplexing, frame relay and ATM.

Thanks to the following people for their invaluable contributions to this project:

Linda Robinson
Moon Kim
Paul Braun
Shawn Walsh
Systems Management and Networking ITSO Center, Raleigh.

Beverly George
Bruce Sinclair
Caglan Aras
Chris Meyer
David Kesselring
Dick Heath
Don Jones
George Viola
Harold Fuller

Jack Lucas
Jim Blye
Ken Young
Ovies Brabson
Steve Bream
Sudarshan Ganapathi
Sushil Jhawar
IBM Research Triangle Park

Peter Krrlander
IBM Sweden

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 193 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Home Pages at the following URLs:

For Internet users	`http://www.redbooks.ibm.com`
For IBM Intranet users	`http://w3.itso.ibm.com/redbooks`

- Send us a note at the following address:

  `redbook@vnet.ibm.com`

Chapter 1. Concepts of Audio and Video Compression

This chapter contains useful information about audio and video signals, describes analog and digital standards, and explains the compression methods used today.

1.1 Introduction

The creation of pictures has been a human activity since the beginning of humanity. Initially this involved painting or carving on rocks. But while rocks are nice and permanent, they are often difficult to transport due to their size. Human beings like to share their pictures, so the search for methods of creating nice pictures on transportable medium was on.

Most of the picture-creating methods have developed into major industries, but an interesting phenomenon has occurred in the past ten years. Most of these technologies are merging toward a more unified approach to imaging, the digital approach. This trend can be attributed to many technological advances, the most important of which is the integrated circuit.

Let′s look briefly at the history of picture technologies.

1.1.1 Motion Pictures

Work in the area of moving pictures from the projection of images using film began in the 1860s. Thomas Edison is often credited with the invention of the first viable motion picture system in 1892, though some people claim he just combined the discoveries of others. His device was called a *kinetograph*. Two major hurdles in the early development of motion pictures were the design of a projection system that eliminated a visible flicker between images, and a lamp with sufficient lighting power to project images onto a distant screen.

Standards have been an issue throughout the evolution of motion pictures. Major characteristics of motion pictures that have been standardized to some degree include the frame rate of projection, the aspect ratio of the film, the type of color-reproduction system used, the type of sound reproduction used, and the synchronization of images to sound.

The *frame rate* of a motion picture is the number of individual images displayed each second during the projection. A variety of frame rates were used for silent films, but the addition of sound to films required the standardization of the frame rate. In 1927, the frame rate for motion pictures was standardized at 24 frames per second. In fact, although only 24 distinct images are shown each second, each image is shown twice so that 48 images are seen each second. This is still the predominant standard used by film motion pictures today.

The *aspect ratio* of an image is the ratio of its width to its height. Before the mid-1950s, the standard aspect ratio for motion pictures was 1.33 to 1 (or 4 to 3). After that time the United States standardized the aspect ratio to 1.85 to 1, while the European countries standardized on 1.66 or 1.75 to 1. Other aspect ratios have been used over the years including the Cinerama system with an aspect ratio 2.75 to 1. Most people feel that a wider picture increases the realism of the motion picture. This perception is due in part to the fact that most action takes place in the horizontal direction. An additional consideration is that the wider screen involves more peripheral vision, which appears to be important to creating a feeling of visual immersion in the scene.

One important issue in motion pictures is the synchronization of the images with sound. One way to encode sound on a film strip is by optical encoding. The decoding and playback of the sound requires a separate mechanism from the projection of the images. Therefore, on a standard 35-mm film strip, optical sound is recorded 21 frames in advance of its corresponding image. On a 16-mm print the sound is recorded 26 frames in advance.

Another method of sound encoding is on magnetic tape. On a standard 70-mm magnetic sound print, the sound is recorded 28 frames behind its corresponding image. On standard 35-mm film print it is 23 frames behind. The synchronization of sound and images is also a problem when the motion picture and sound data are in digital form.

The hope of many film professionals is that some day motion pictures can be produced without film in a digital data format. Currently, film can capture higher resolution, higher-quality images than digital cameras. Film will continue to be the predominant medium for motion pictures until digital cameras can achieve comparable resolutions (at comparable prices). In the meantime, systems have been developed that take film images, transfer them to digital form for editing and incorporation of special effects, and then transfer them back to film for display.

1.1.2 Television

The concept of sending a moving picture over an electrical wire dates back to the 1870s. The initial ideas envisioned a system much like a fax machine that could send an image decomposed into a group of dots, but every dot in the image would be communicated simultaneously. This idea quickly gave way to the notion of sending the dots of each image in a serial fashion and reconstructing the image after it was received.

The first patent of a complete television system was issued to Paul Nipkow in Germany in 1884. It used a mechanical rotating drum to scan an image, both for sending and receiving. Mechanical systems dominated the research in this area for many years. In 1926 J.L. Baird, a developer in England, gave the first demonstration of true television. His system had a resolution of 30 lines and it updated the screen approximately 10 frames every second. It was crude, but it proved the feasibility of the idea and stimulated further research.

Mechanical systems lacked sensitivity, making it difficult to extend them to higher resolution images. It was recognized by most developers that a good-quality image of reasonable size required at least 100,000 elements and preferably 200,000 elements, assuming a screen approximately 12 inches high was viewed from a distance of 5 to 8 feet. This number can be calculated from the resolving power of the human visual system. For a square image, 200,000 elements requires approximately 447 scan elements per line. This was beyond the feasible limits of mechanical systems.

The idea of a totally electrical system was first proposed by a Scottish electrical engineer, A. Campbell Swinton, in 1908. His idea was to use magnetically deflected beams in CRTs to scan an image. The idea was too advanced for the technology of 1908, but it was implemented by V.K. Zorykin′s iconoscope camera tube, which was patented in 1923. Based on these developments, an all-electronic television was demonstrated by the Radio Corporation of America (RCA) in 1932. It initially used 120 scan lines which was rapidly increased to 343.

The Electric and Musical Industries (EMI) of Great Britain began television research in 1931 under the direction of Isaac Shoenberg. Their team developed a complete and practical system by 1935 and launched the first public television service in 1936. Shoenberg was very concerned with the standardization of the image signal. He proposed an image signal using 405 scan lines updated at a rate of 25 times per second. To decrease the visible flickering effect that can occur at this frame rate, the scan lines were *interlaced*. The screen is divided into even- and odd-numbered scan lines. All of the odd scan lines for the screen are updated first in 1/50th of a second; then the even scan lines are updated. The net effect is a completely new image displayed 25 times per second. The frame rate of 50 updates per second was chosen to match the electrical power line frequency used in Great Britain (50 Hz). If other update frequencies were used, unwanted distortions in the picture (artifact) would occur due to competing eletromagnetic fields. The EMI standard formed the basis for British television up until 1964.

The United States began regular television broadcasts in 1941. The United States standardized on a signal that included 525 scan lines per picture updated 60 times per second interlaced, resulting in a total update of the screen 30 times per second. The actual refresh rate is 59.94 fields per second, but its exact value is not significant for most television viewers. The update rate of 60 times per second was based on the electrical power line frequency used in North America (60 Hz). Meanwhile, European countries other than England standardized on a signal of 625 lines at 25 frames per second. Most other countries in the world began their television services in the 1950s and chose either the United States standard or the European standard based on the frequency of alternating current their electrical power systems generated.

The idea of color television existed almost from the very beginning of television research, but color was a more difficult problem.

The human retina has three types of color photoreceptor cone cells, which respond to incident radiation with different response curves. There is a fourth type of photoreceptor cell also present in the retina, but it is only effective at extremely low light levels, and although important for vision, it plays no role in color image reproduction.

Because there are exactly three types of color photoreceptors in the eyes, three numerical components are necessary and sufficient to describe a color. All colors are perceived as a combination of the so-called primary colors: red, green and blue. However, the colors produced by reflective systems such as photography, printing or paint are not only a function of the colorants but also the ambient illumination.

Using these concepts, the National Television Standards Committee (NTSC) defined the transmission of video signals in a luminance (or brightness) and chrominance (or hue) format, rather than a format involving the three color components of the television phosphors. Other television standards used worldwide, such as PAL (Phase Alternation by Line) and SECAM (Systeme Electronique Couleur Avec Memoire), are based on the same three components. They are defined as:

- Luminance: Y

- Red Chrominance: Cr = Y - Red
- Blue Chrominance: Cb = Y - Blue

This representation enables backward compatibility with black and white TVs, which use the luminance (Y) component to build the images. Another reason for choosing these components is that the red chrominance (Cr) consists of a hue that is similar to human flesh tones.

(NTSC is sometimes jokingly referred to as ″Never Twice the Same Color″. SECAM also has its non-official version of ″Something Essentially Contrary to the American Method″.)

1.1.3 HDTV Developments

The current focus of research is on high-definition TV (HDTV). The main emphasis of HDTV is to provide a better-quality image by doubling the resolution in both the horizontal and vertical directions. This requires the transmission of four times as much information per frame. In addition, there is a push to modify the aspect ratio of the image to increase the width of the image area. Several aspect ratios have been proposed, the most common of which is 16 to 9 (1.78 to 1).

The Federal Communications Committee (FCC) in the United States is currently working on standardizing a new HDTV signal to broadcast in the United States. The FCC has stated that the new HDTV signal must fit within the same signal bandwidth that the current NTSC signal does (approximately 6 MHz). The committee focused its study on six proposals, one of which was expected to be chosen as the new standard. Of the six proposals, five used digital encoding. During the evaluation process, the submitting organizations and consortia finally agreed to create a ″grand alliance″ which would combine the best parts of all of the digital approaches. It is thus almost certain that the next standard for TV broadcast will be a digital encoding of raster data. This is a significant development in the evolution of TV and will have a major impact on many other industries. Although Europe had an anglo-based system (HDMac) in advanced development, and Japan has demonstrated an analog-based system; both of these regions are now looking toward digital-based systems.

1.2 Analog Video Formats

The analog video formats that are in use throughout the world were defined by the television standard bodies.

Table 1 shows the transmission and resolution characteristics of these different video formats.

Table 1. Analog Baseband Video Formats

Description	NTSC	PAL	SECAM
Video Resolution (lines x columns)	525 x 720	625 x 720	625 x 720
Frames per Second	29.97	25	25
Where Used	US, Japan	Most of the World	France, Russia
Effective Resolution	640 x 480	768 x 576	768 x 576

Video signals contain a blanked portion that is used for synchronization purposes, but are not displayed on the screen. The effective resolution is measured after extracting this non-active video used for synchronization.

This non-active portion of the video signal is called Vertical Blanking Interval (VBI) and consists of 24 lines (12 per field). The VBI can be used to transmit ancillary data such as closed captioning, extended data services, video test patterns for ghost cancellation, editing time codes, time-of-day information, emergency broadcasts and data transmission services.

The baseband video signal is transmitted separately from the audio signals. Therefore, at least three cables are required to connect a VCR to a camera or to a TV set: video, audio-left and audio-right.

There are different ways to store and transmit baseband video:

- **Composite Video:** In this case, only one signal is used to store and transmit luminance and chrominance. It has the advantage of requiring less bandwidth and only one cable, but due to cross-talk, the video quality is reduced. This is available from your VCR.
- **S-Video:** The luminance and chrominance signals are transmitted separately, offering a better video quality, with crisper details. This is usually available on high-end equipment, such as Hi8 cameras.
- **Component Video - RGB:** This is the most sophisticated way to transmit video signals. The three video signal components are sent separately and maximum resolution is achieved. Three cables are required. This is found only in professional equipment.

TV signals are modulated signals. This means that the baseband audio and video signals are frequency-modulated and are transmitted in specific frequencies, according to the associated TV channel. Audio in the TV signal is carried as a stereo FM transmission in NTSC and PAL broadcasts.

Your TV antenna receives all the channels that are modulated in that frequency range and then the TV tuner selects the desired channel among all those transmitted.

1.3 Digital Audio and Video Formats

The following sections describe the encoding standards adopted for digitization of analog audio and video signals. This process is required before compression can be applied and plays an important role in determining the final signal quality.

1.3.1 Digital Video Encoding

Video digitization is the process of converting the luminance and chrominance information present in the analog signal into digital format. There is no compression involved at this point.

The former Comite Consultatif International des Radiocommunications (CCIR), now the International Telecommunications Union - Radiocommunications (ITU-R), developed several recommendations for transmission and storage of video signals. The Standard Digital Video Format was defined in the CCIR-601 and is also called *D-1 format.* This format was developed for digital video tape and determines how a video signal must be digitized.

The process of video encoding is represented in Figure 1 and shows a D-1 signal created after an analog video input. High-quality video compression always starts with a D-1 encoded digital video signal.

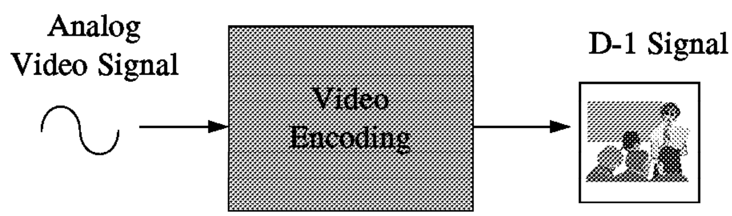

Figure 1. Block Diagram of Video Encoding

Let's understand a little more about the terms used to define these video formats.

The word *pixel* stands for picture element and is used to represent the color characteristics of each element of the picture. Pixels can carry more or fewer details of information depending on the profile chosen:

- Profile 4:2:2 determines that each pixel uses 8 bits to represent luminance information (Y), 4 bits for red chrominance (Cr) and 4 bits for blue chrominance (Cb). This amounts to 16 bits per pixel.
- Profile 4:2:0 allocates 8 bits for luminance, 2 bits for Cr and 2 bits for Cb. In this case a pixel is made of a total of 12 bits.

The CCIR-601 standard defines the D-1 format in the following way:

- 720 pixels/line x 480 lines/frame x 29.97 frames/second for NTSC
- 720 pixels/line x 576 lines/frame x 25 frames/second for PAL

Note that only active lines are considered. Inactive lines are not always digitized.

Figure 2 on page 7 illustrates a video signal frame.

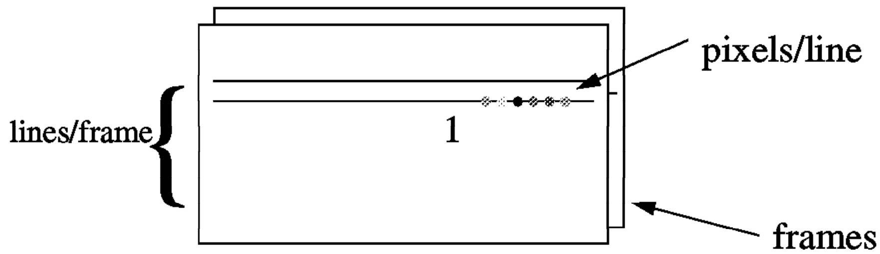

Figure 2. Contents of a Video Frame

In both cases, the total amount of bandwidth needed to store or transmit a D-1 digital video is roughly 166 Mbps (for 4:2:2 profile) or 124 Mbps (for 4:2:0 profile).

This means that the transmission of a D-1 digital video stream would require 74.6 GB/hour. These values show how important video compression is to make digital video transmission possible in practical terms.

In order to be able to store a reasonable amount of video on a 650 MB CD-ROM, originally intended for playback at single speed, video resolution has to be reduced to a quarter of the size of the CCIR-601 D-1 standard to a VHS quality called *SIF* (Source Input Format). SIF pictures can easily be derived from a CCIR-601 frame using filtering and subsampling. The resulting bit rate for the television standards are:

- 360 pixels/line x 240 lines/frame x 29.97 frames/second for NTSC
- 360 pixels/line x 288 lines/frame x 25 frames/second for PAL

The bandwidth needed for transmitting a TV video signal using SIF format drops to approximately 31 Mbps.

In practice, the coders only use multiples of 16 bits as the minimum coded unit, also called a macroblock. Since the horizontal resolution of 360 pixels in a SIF picture is not divisible by 16, the four leftmost pixels and the four rightmost pixels of each line are discarded. That is why you may see specifications for the digital video formats listing what is called the significant pixel area. These values are 704 pixels/line for the CCIR-601 and 352 pixels/line for the SIF format.

Other digital video formats have been defined to match the computer displays. Starting with the square pixel *VGA*, which is a 640x480 matrix, the *CIF* (Common Intermediate Format) was defined, with one quarter of the VGA resolution, or 320x240. By further dividing the CIF format by a factor of four, we obtain the *QCIF*, which stands for Quarter-CIF and is used for applications with lower video quality.

Higher resolution formats have also been created from the VGA format. They are called *Super-VGA*, or *SVGA*.

Several digital video formats are shown in Figure 3 on page 8.

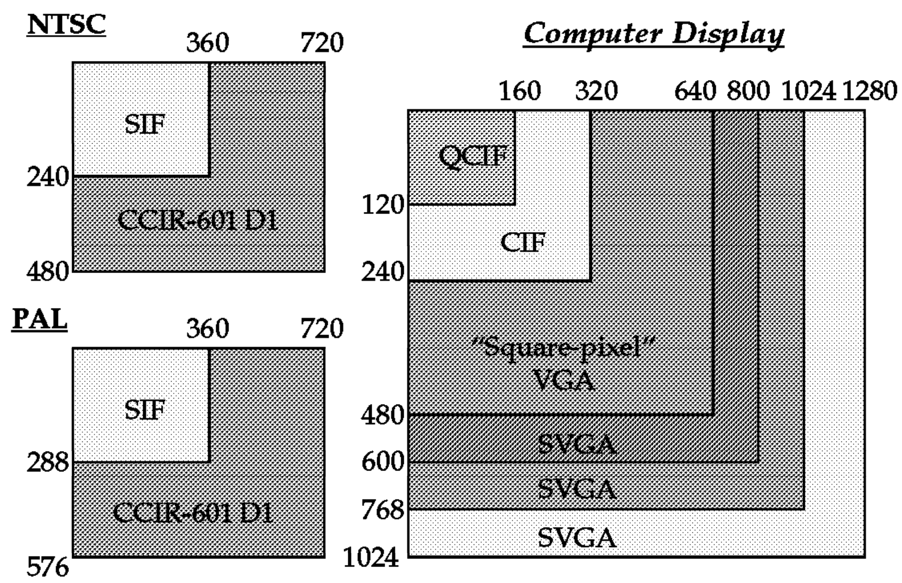

Figure 3. Digital Video Formats

1.3.2 Digital Audio Encoding

Audio is digitized at a sampling rate, using a selected sampling size. Of course, different levels of quality are achieved based on the number of samples per second and the number of bits used to represent each sample.

The basic rule for audio digitization is based on the Nyquist Theorem, which determines that the sampling rate must be greater than two times the highest frequency of the analog signal.

Here are examples of some systems and their corresponding frequency ranges:

- The human perception: 20 Hz to 22 KHz
- FM broadcasts: 70 Hz to 15 KHz
- Telephone speech: 1 KHz to 3.5 KHz

That is why CD recordings use a 44.1 KHz sampling frequency, approximately twice the maximum frequency response of the human auditory system. In digital telephony, where high audio fidelity is not as important, signals are sampled at 8 KHz. Some of the commonly used encoding formats are illustrated in Table 2.

Table 2. Digital Audio Quality

Quality	**Samples/sec**	**Bits/sample**	**Bandwidth required**
Telephone	8000	8	64 kbps
CD Digital Audio	44100	16 x 2 channels	1.4 Mbps
HDTV, DAT	48000	18 x 6 channels	5.1 Mbps

1.4 Video Compression

In recent years, there have been significant advancements in algorithms and architectures for the processing of audio and video signals. One of the exciting prospects of such progress is that multimedia information comprising image, video and audio, has the potential to become just another data type. This usually implies that multimedia information will be digitally encoded so that it can be manipulated, stored and transmitted along with other digital types. The digitization process, covered in the previous section, produces a digital video signal, such as a D-1 signal.

Compression is a process intended to yield a compact digital representation of a signal and to minimize its bit rate. Figure 4 illustrates this process. There are many applications that benefit when audio and video signals are available in compressed form. Without compression, most of these applications would not be feasible. Standard data encoding methods among systems and applications are also essential to promote interoperability.

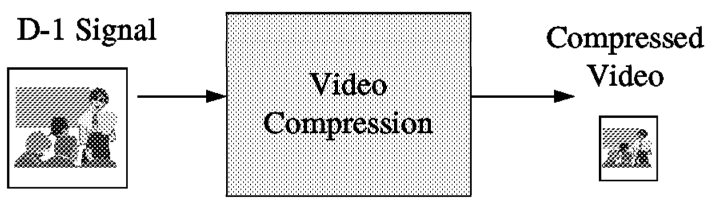

Figure 4. Block Diagram of Video Compression

Two important standardization efforts were started in the late 1980s. One is the ITU-T standard for videoconferencing and video-telephony, known as H.261. The other one came under the name of MPEG (Moving Pictures Experts Group) from ISO/IEC in order to define a video coding algorithm for applications on digital storage media. When the MPEG committee began the task of specifying a syntax for compressed digital video, its goal was to deliver video on a compact disc, taking into account its low data transfer rate of 1.416 Mbps.

Being aware that it was impossible to represent a CCIR 601 (D-1 format) resolution image at such a low data rate, the committee specified the one-quarter resolution image (SIF format) as the standard input format. When decoded, the SIF video signal is expanded to fill a full television screen, resulting in an image quality similar to VHS tape. In addition, audio coding was added and the scope of the targeted applications extended to cover most applications, from multimedia systems to Video-on-Demand. The activities of JPEG (Joint Photographic Experts Group) played an important role in the definition of MPEG.

Although JPEG was originally intended for still-image compression, it can also be applied for motion pictures, considering the fact that a video sequence is nothing more than a sequence of still images. These still images can be coded individually and displayed sequentially using JPEG, but the coded sequence will not take into consideration frame-to-frame redundancies that give an additional compression factor.

MPEG exploits the temporal redundancy present in any video sequence in order to maximize the compression ratio. MPEG′s first effort led to the MPEG-1 standard that was published in 1993 as ISO/IEC 11172. It is divided in three parts: audio compression, video compression and system level multiplexing for applications that need video and audio to be played back in close synchronization. MPEG-1 is being used in a variety of applications. CD-I and Video-CD technology use MPEG-1 as the compression algorithm for video and audio. MPEG-1 was designed to support video coding up to 1.5 Mbps with VHS quality, audio coding at 192 kbps/channel (stereo CD-quality), and is optimized for non-interlaced video signals.

Broadcast television equipment makers immediately recognized the potential of MPEG technology to increase the channel efficiency of satellite transponders and cable networks, but the broadcast industry was not limited to present-day compact disc bandwidths and was unwilling to settle for VHS resolution. Consequently, the MPEG committee developed a second standard, especially designed to represent CCIR-601 resolution video (D-1 format) at data rates of 4.0 to 15.0 Mbps.

MPEG′s second effort started in 1990. The main objective was to design a compression standard capable of different qualities depending on the bit rate, from TV broadcast to studio quality. That work led to the MPEG-2 standard, which is based on MPEG-1 but is more sophisticated and optimized for interlaced pictures. The MPEG-2 standard is capable of coding video signals that range from standard TV to HDTV (High-Definition TV). In the audio part of the standard, it supports multichannel surround sound coding while being backwards compatible with MPEG-1 encoded signals.

1.4.1 The MPEG Standards

Each of the two standards, MPEG-1 and MPEG-2, is divided into three sections: systems, video and audio.

The systems specification addresses the combination of separate audio and video streams into a single stream for storage or transmission. It also provides a mechanism for synchronizing audio and video by MPEG decoders. The following discusses both MPEG-1 and MPEG-2 system specifications.

1.4.1.1 The MPEG-1 Stream Structure

In its most general form an MPEG-1 stream is made up of two layers. The system layer contains timing and other information needed to demultiplex the audio and video streams and to synchronize audio and video during playback. The compression layer includes the compressed audio and video.

The system demux extracts the timing information from the MPEG stream and sends it to the other system components. The system demux also demultiplexes the video and audio streams and sends each to the appropriate decoder. The video decoder decompresses the video stream while the audio decoder decompresses the audio stream.

1.4.1.2 MPEG-2 Streams

MPEG-2 defines two kinds of system streams: program and transport. Both are multiplexed streams consisting of video and audio elementary streams and both are subdivided into packets for transmission;

- **Program Stream**

  A single stream resulting from the combination of one or more elementary streams, that share a common time base. The MPEG-2 program stream is similar to the MPEG-1 system stream. The program stream is designed to be used in relatively error-free environments such as multimedia applications. The packets in a program stream can be of any length.

- **Transport Stream**

  A single stream, resulting from the combination of one or more programs, in which a program is a collection of elementary streams with a common time base. The transport stream is designed for relatively error-prone environments such as broadcast applications. Transport stream packets are fixed at 188 bytes in length. Figure 5 illustrates how the transport streams are formed.

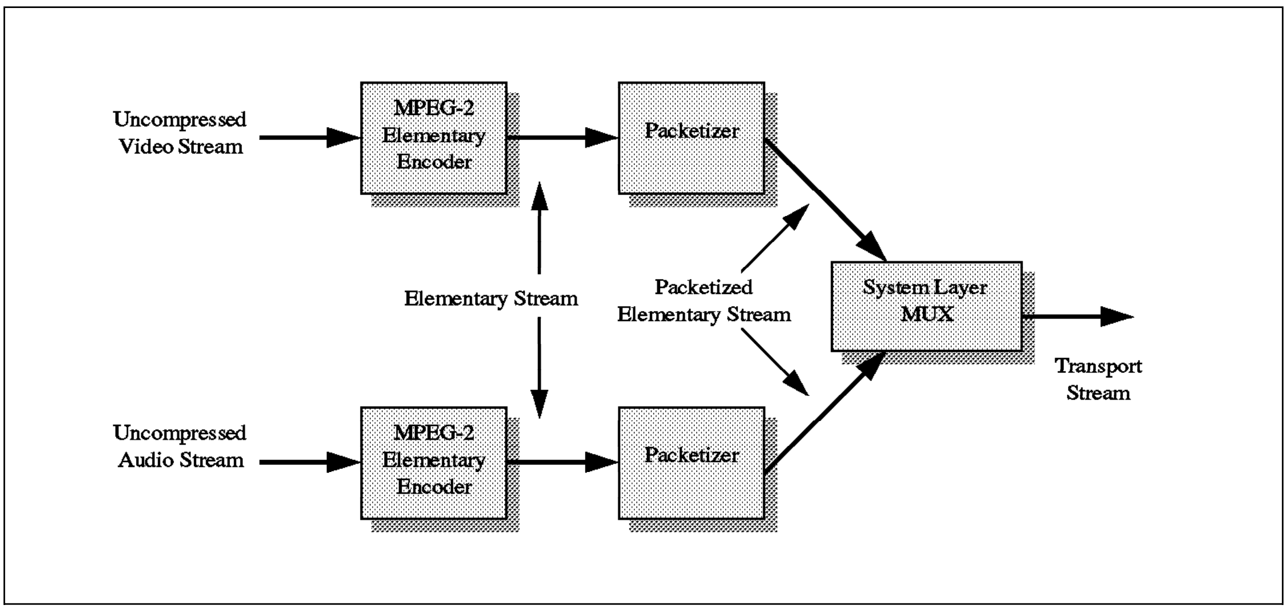

Figure 5. Simplified Block Diagram of Transport Stream

1.4.1.3 Video Stream Data Hierarchy

The MPEG-1 and MPEG-2 standards define a hierarchy of data structures in the video stream as listed below:

- **Video Sequence**

  Consists of a sequence header, one or more groups of pictures and an end-of-sequence code. The video sequence is another term for a video stream as defined above.

- **Group of Pictures (GOP)**

  A series of one or more pictures intended to allow random access into the sequence.

- **Picture**

  The primary coding unit of a video sequence. A picture consists of three rectangular matrixes representing luminance (Y) and two chrominance (CbCr) values. The Y matrix has an even number of rows and columns. The Cb and Cr matrices are one-half the size of the Y matrix in each direction (horizontal and vertical).

- **Slice**

  One or more contiguous macroblocks. The order of the macroblocks within a slice is from left to right and top to bottom. Slices are important in the handling of errors. If the bit stream contains an error, the decoder can skip to the start of the next slice. Having more slices in the bit stream allows better error concealment but uses bits that could otherwise be used to improve picture quality.

- **Macroblock**

  The basic coding unit in the MPEG algorithm. It is a 16x16 pixel segment in a frame. Since each chrominance component has one-half the vertical and horizontal resolution of the luminance component, a macroblock consists of four Y blocks, one Cr block, and one Cb block.

- **Block**

  The smallest coding unit in the MPEG algorithm. It consists of 8x8 pixels and can be one of three types: luminance (Y), red chrominance (Cr), or blue chrominance (Cb). The block is the basic unit in intraframe coding.

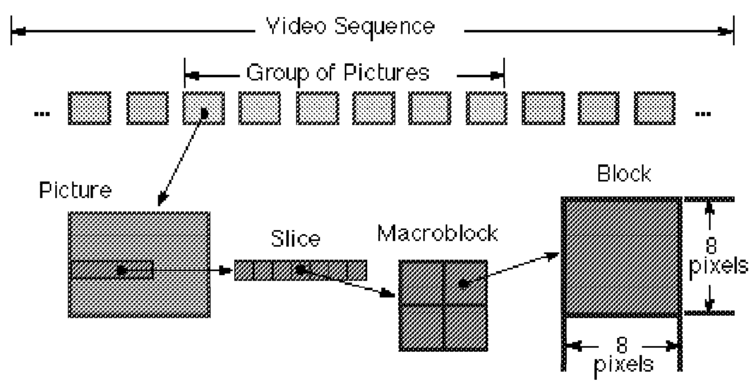

Figure 6. MPEG-2 Video Sequence

1.4.1.4 Information Redundancy

Much of the information in a picture within a video sequence is similar to information in a previous or subsequent picture. The MPEG-1 and MPEG-2 standards take advantage of this redundancy to represent some pictures in terms of their differences from a reference picture.

These are the redundancy types that are exploited by MPEG-2:

- **Spatial Redundancy**

Large areas of similar information within a frame. Figure 7 on page 13 shows an example of a large rectangle that appears in a specific frame. This area can be compressed independently of any other frame.

- **Temporal Redundancy**

  Similar information in different frames. In the example shown in Figure 7 the teacher had her position changed from one frame to the other, but her shape remained the same. This kind of information can also be used for compression.

MPEG-2 uses Discrete Cosine Transform (DCT) and entropy encoding to deal with spatial redundancies (intraframe coding), and motion compensation and motion estimation for the temporal redundancies (interframe coding).

Huffman encoding is also used to efficiently represent repeated data.

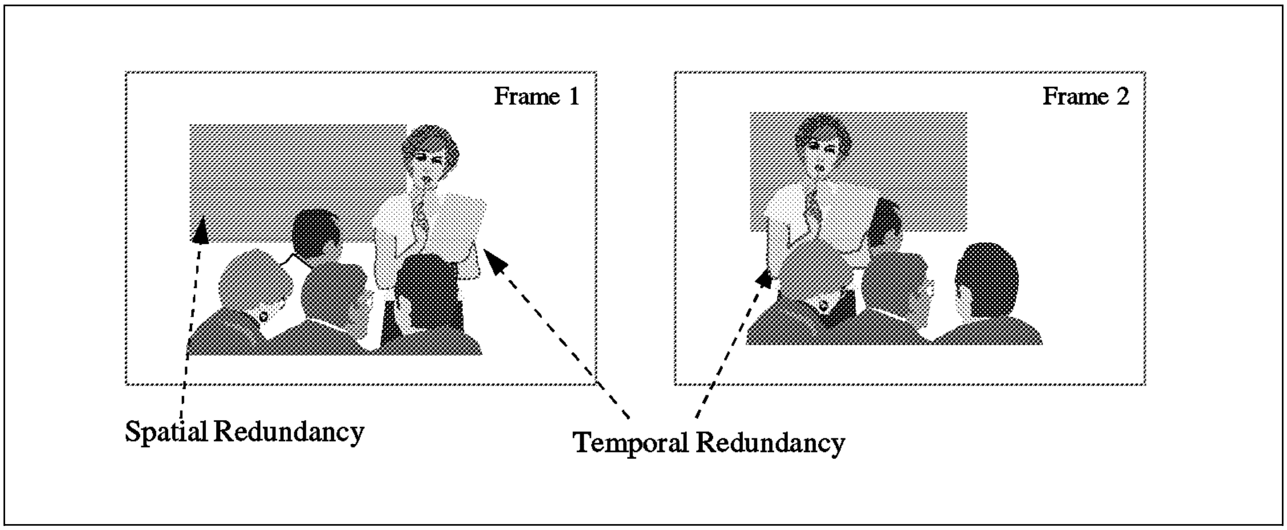

Figure 7. Spatial and Temporal Redundancy

1.4.1.5 Picture Types

The MPEG standard specifically defines three types of pictures:

1. Intra pictures (I-pictures)
2. Predicted pictures (P-pictures)
3. Bidirectional pictures (B-pictures)

These three types of pictures are combined to form a group of pictures, as shown in Figure 8 on page 14.

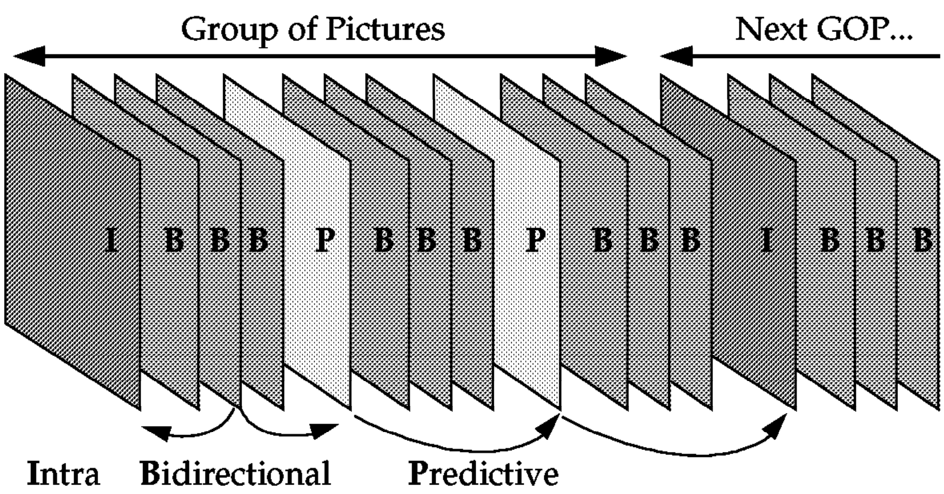

Figure 8. Group Of Pictures (GOP)

- **Intra Pictures**

  Intra pictures, or I-pictures, are coded using only information present in the picture itself and do not depend on another picture to reconstruct the image. They only exploit spatial redundancy to compress information within the frame and provide moderate compression. Typically they use about 2 bits per coded pixel.

- **Predicted Pictures**

  Predicted pictures, or P-pictures, are coded with respect to the nearest previous I- or P-pictures and exploit temporal and spatial redundancy to compress a video frame. An anchor frame must be referenced in order to reconstruct the image, and the references are always made to past frames. This technique is called forward prediction. Like I-pictures, P-pictures can also serve as a prediction reference for B-pictures and future P-pictures. Moreover, P-pictures use motion compensation to provide more compression than is possible with I-pictures.

- **Bidirectional Pictures**

  Bidirectional pictures, or B-pictures, are pictures that use both a past and future picture as a reference. This technique is called bidirectional prediction. B-pictures provide the most compression, since they use the past and future pictures as a reference. However, the computation time is the largest.

As an example, Figure 9 on page 15 shows how the pictures can be scanned searching for similar blocks in different frames. The search space defines the extent of the search to find a similar block. The greater the search space the more likely a similar block is found, but more processing and memory is needed.

Predicted pictures (P) only search for blocks in past frames while Bidirectional pictures (B) search past and future frames.

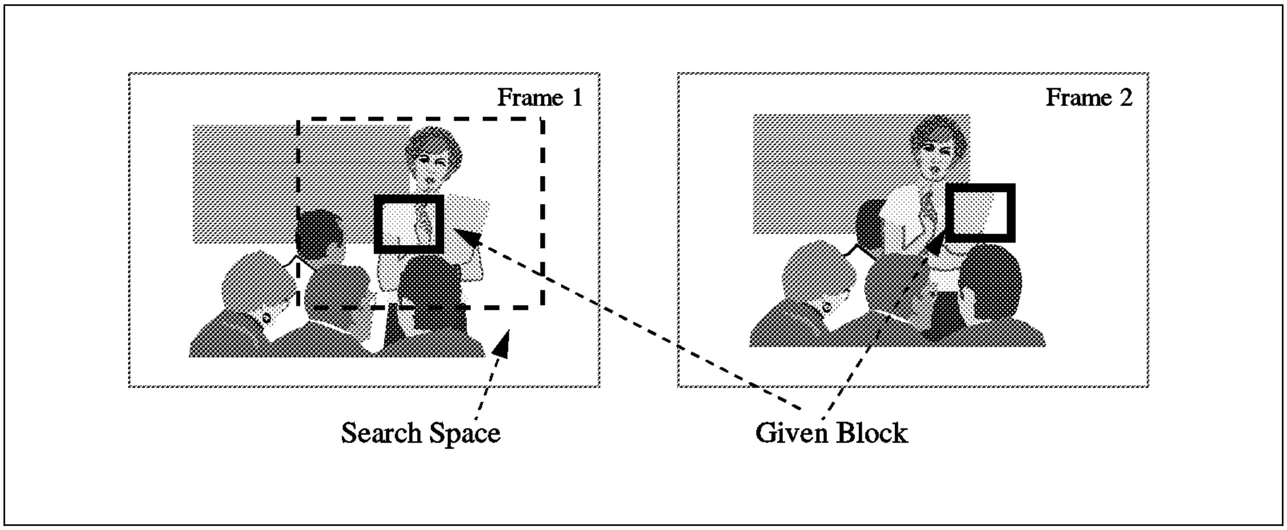

Figure 9. Search Space

1.4.1.6 Synchronization

The MPEG standards provide a timing mechanism that ensures synchronization of audio and video.

The MPEG-1 standard defines two parameters used by the decoder: the system clock reference (SCR) and the presentation timestamp (PTS).

The MPEG-2 standard adds the program clock reference (PCR) and also provides for the SCR and PCR to have extensions with a resolution of 27 MHz.

- **System Clock Reference**

  An SCR is a snapshot of the encoder system clock. The SCRs used by the audio and video decoder must have approximately the same value. To keep their values in agreement, SCRs are inserted into the MPEG stream as often as every 0.7 seconds (minimum) by the MPEG encoder, then extracted by the system decoder and sent to the audio and video decoders. The video and audio decoders update their internal clocks using the SCR value sent by the system decoder.

- **Presentation Timestamps**

  PTSs are samples of the encoder system clock that are associated with some video or audio presentation units. A presentation unit is a decoded video picture or a decoded audio time sequence. The encoder inserts PTSs into the MPEG stream as often as every 0.7 seconds (minimum). The PTS represents the time that the video picture is to be displayed or the starting playback time for the audio sequence. The video decoder either deletes or repeats pictures to ensure that the PTS matches the current value of the SCR when a picture with a PTS is displayed. If the PTS is earlier (has a smaller value) than the current SCR, the video decoder discards the picture. If the PTS is later (has a larger value) than the current SCR, the video decoder repeats the display of the picture.

- **Program Clock Reference**

  PCRs are used only in MPEG-2. The PCR is used in the transport stream in the same way that the SCR is used in an MPEG-1 system stream. Since each program can have its own time base, a transport stream containing multiple programs has a separate PCR field for each program.

1.4.1.7 Interlaced Video and Picture Structures

MPEG-2 supports two scanning methods: one is progressive scanning and the other is interlaced scanning. Interlaced scanning scans odd lines of a frame as one field (odd field), and even lines as another field (even field). Progressive scanning scans the consecutive lines in sequential order.

An interlaced video sequence uses one of two picture structures: frame structure and field structure. In the frame structure, lines of two fields alternate and the two fields are coded together as a frame. One picture header is used for two fields. In the field structure, the two fields of a frame may be coded independently of each other, and the odd field is followed by the even field. Each of the two fields has its picture header.

The interlaced video sequence can switch between frame structures and field structures on a picture-by-picture basis. On the other hand, each picture in a progressive video sequence is a frame picture.

1.4.1.8 Profiles and Levels

MPEG-2 is designed to support a wide range of applications and services of varying bit rate, resolution, and quality. MPEG-2 standards define four profiles and four levels for ensuring interoperability of these applications.

Table 3 and Table 4 summarize the characteristics of these levels and profiles.

Table 3. MPEG-2 Levels

Level	Max. Resolution x fps	Pixels/sec	Max. Bit Rate	Application Example
Low	352 x 288 x 30	3.05 M	4 Mbps	SIF, VHS
Main	720 x 576 x 30	10.40 M	15 Mbps	CCIR-601, studio TV
High-1440	1440 x 1152 x 30	47.00 M	60 Mbps	4xCCIR-601, HDTV
High	1920 x 1152 x 30	62.70 M	80 Mbps	SMPTE 240, video production system

Table 4. MPEG-2 Profiles

Profile	Comments
Simple	Encoding of I and P pictures. Uses 4:2:0 macroblocks.
Main	Encoding of I, P and B pictures. Uses 4:2:0 macroblocks.
Main+	Equivalent to Main, with Spatial and SNR scalability.
Next	Equivalent to Main+ with 4:2:2 macroblocks.

1.4.1.9 Comparing MPEG-1 with MPEG-2

The MPEG group defined two compression algorithms: MPEG-1 and MPEG-2. Each algorithm has been specifically targeted at different bit rates, with MPEG-2 at higher rates than MPEG-1.

The essential difference between MPEG-2 and MPEG-1 video is the incorporation of field-based motion prediction. Figure 10 shows the supported video quality for different transmission rates.

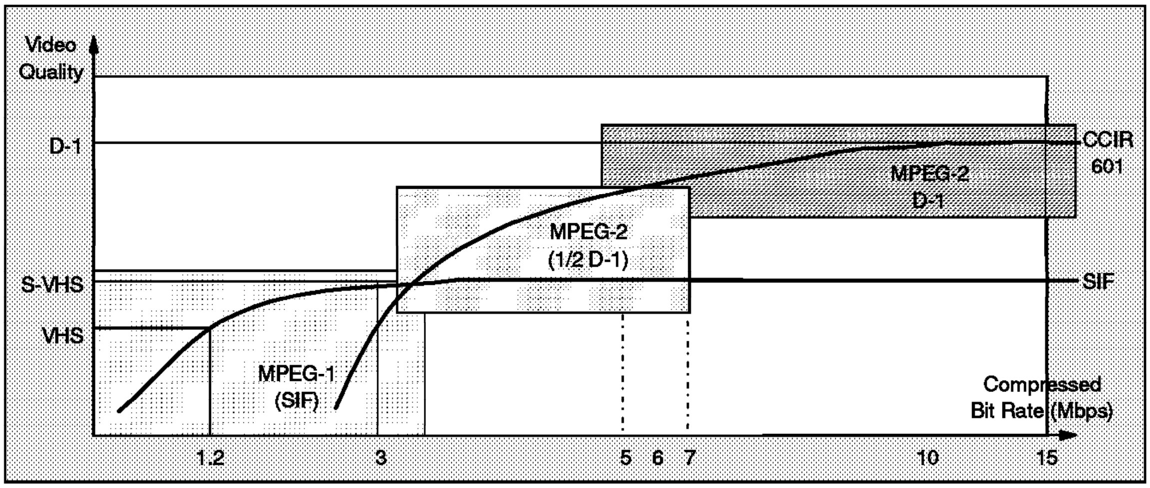

Figure 10. Transmission Rate vs Video Quality for MPEG Standards

The CCIR-601 curve corresponds to compressing the video at the full input resolution (D-1 format: 720 x 576 for PAL). As the compressed bit rate is reduced for this curve, the output video quality is also gradually reduced until somewhere below 5 Mbps, when the quality drops off very steeply.

On the other hand the curve that represents the SIF format is unable to provide equal quality at the highest rates, but when the rate drops below about 3.5 Mbps, it outperforms the CCIR-601 curve. The SIF is a reduced resolution format, as explained in the former section (360 x 288 for PAL).

Although MPEG-1 can be run at high bit rates and at full CCIR-601 resolution, it always processes frames and not fields. This limits the attainable quality even at high bit rates (above 5 Mbps) and pushes the development of the MPEG-2 algorithm that can process individual fields. However, to take advantage of this capability and produce better pictures, the MPEG-2 algorithm has to be run at the higher rates and with interlaced input. When MPEG-2 is run at the lower bit rates and at reduced SIF resolution, the significant additional cost and complexity of an MPEG-2 encoder is wasted, since it does not produce better quality than an MPEG-1 SIF encoder.

The graph shows that the optimal choices are:

- Below 3.5 Mbps - MPEG-1 at SIF Resolution
- Above 5 Mbps - MPEG-2 at CCIR-601 Resolution

1.5 Audio Compression

Digital audio is an integral part of any video or multimedia application and may consume a considerable portion of the overall bandwidth.

Assuming a digital audio signal is generated as described in 1.3.2, "Digital Audio Encoding" on page 8, the bit rate from a stereo CD is close to 1.4 Mbps, and a 10-minute stereo audio clip requires approximately 105 MB of storage. This indicates that efficient storage and transmission of audio requires some form of data compression. One of the main compression techniques for high-fidelity audio is the MPEG audio standard.

The MPEG audio compression algorithm is part of the MPEG-1 standard that addresses not only the compression of video and audio but also their synchronization. It can be applied to streams that combine both video and audio or to audio-only streams. This technique can achieve CD-like quality at 128 kbps per audio channel (5.5:1 compression), or down to 32 kbps for a single channel.

Presently, within ITU, a videoconferencing standard referred to as H.324 is being developed for videoconferencing at bit rates lower than 64 kbps. A speech coding method referred to as G.723 has been proposed as part of the H.324 standard suite. The 6.3 kbps mode of the G.723 provides speech quality equivalent to that provided by the 32 kbps G.726 standard.

With MPEG, audio is compressed by eliminating silences, similarities between the two stereo channels, and the reduction of audio fidelity.

Psycho-acoustic models are used to detect sound patterns that are normally masked by other sounds. These are eliminated during compression.

Audio compression defines a certain bit rate for the compressed material. Typical rates are 96, 192, 256 and 384 kbps. Rates of 256 kbps and above provide CD quality.

The MPEG-2 audio compression standard is an extension of MPEG-1 with the following added features:

- Multichannel Input

  MPEG-2 supports up to five high-fidelity audio channels plus a low-frequency enhancement channel. This is suitable for the compression of audio in HDTV or digital movies.

- Lower Bit Rates

  Compressed bit rates down to 8 kbps are now supported.

- Multilingual Audio Support

  Up to eight commentary channels are supported.

- Additional Sampling Rates

  In addition to the original sampling rates of 32, 44.1 and 48 KHz, MPEG-2 accommodates sampling rates of 16, 22.05 and 24 KHz.

MPEG-2 audio decoders can decode MPEG-1 audio bit streams, and MPEG-1 audio decoders can decode the two main channels of the MPEG-2 audio bit streams.

1.6 Comparison of Compression Methods

Many audio visual compression methods are used in industry today. Some of them have been in use for a long time and others are now being created to support the recent technologies and applications. The various compression schemes address specific quality levels and transmission rate ranges.

They are usually classified as high quality (such as MPEG-2, Motion-JPEG and AD-PCM) or medium quality (such as MPEG-1). We briefly describe the most commonly used techniques as well as the ITU communications standards.

1.6.1 AD-PCM

This is a fairly old technology. AD-PCM stands for Adaptive Differential Pulse Code Modulation, and basically consists of an adaptive digital sampling of the video signal.

The following list describes the characteristics for AD-PCM:

- Standardization

  It is not standard, so equipment from different vendors cannot interoperate.

- Video Compression

  It achieves a 3:1 compression rate, resulting in a fixed 36 Mbps data stream.

- Audio Compression

  Audio is carried uncompressed at 1.5 Mbps.

1.6.2 JPEG/M-JPEG

JPEG was developed in the late 1980s for still images but can be used for video considering that it consists of a sequence of still pictures. JPEG codes the still images individually and plays them sequentially, but it does not consider inter-picture redundancies. The audio is transmitted separately from the video and they are combined at the destination through proprietary schemes. M-JPEG or Motion-JPEG adds nonstandard extensions to JPEG to make it more suitable for motion video transmission.

The following list describes the characteristics for JPEG/M-JPEG:

- Standardization

  JPEG is standardized only for still imagery, as ITU T.81.

  M-JPEG is not standard.

- Video Compression

  Only spatial redundancy (I-pictures) are exploited. It was never designed for motion video compression, but is useful in video editing.

  Typical compressed rates are between 15 Mbps and 30 Mbps.

- Audio Compression

  There is no standard audio layer.

- System Layer

  There is no standard system layer.

Usually video and audio streams are transmitted separately with proprietary format time stamps and are combined at the destination.

The lack of system layer prevents the exchange of encoded material.

1.6.3 MPEG-1

MPEG-1 was developed in the early 1990s to support video coding up to 1.5 Mbps with VHS quality and audio coding with stereo CD quality. MPEG-1 and MPEG-2 are described in detail in the previous sections.

The following list describes the characteristics for MPEG-1:

- Standardization

  MPEG-1 is the ISO/IEC 11172 standard, but is not an ITU standard.

- Video Compression

  Its video layer is defined in ISO 11172-2.

  Temporal and spatial redundancies are exploited with I, P and B pictures.

  It is usually limited to SIF frame size.

  Video transmission rates range from 0.8 to 4 Mbps (1.5 Mbps typically).

- Audio Compression

  Its audio layer is defined in ISO 11172-3.

  Audio compression provides CD-quality stereo at 192 kbps per channel with support for two channels.

- System Layer

  Its system layer is defined in ISO 11172-1 and defines a system to store audiovisual content.

  It provides audio and video synchronization through time stamps. It is similar to MPEG-2 program streams.

1.6.4 MPEG-2

MPEG-2 was developed with the purpose of supporting different quality levels at different transmission rates, from TV broadcast to studio quality. MPEG-1 and MPEG-2 are described in detail in the previous sections.

The following list describes the characteristics for MPEG-2:

- Standardization

  MPEG-2 was ratified in 1994 by ISO (ISO 13818) and was merged into the ITU standards.

- Video Compression

  Its video layer is defined in ISO 13818-2 and ITU H.262.

  It exploits temporal and spatial redundancies with I, P and B pictures.

  It can compress frame sizes from SIF, D-1 to 1900x1000.

  Video transmission rates range from 3 to 15 Mbps, with support of up to 50 Mbps.

- Audio Compression

  Its audio layer is defined in ISO 13818-3.

Audio compression provides up to five channel CD quality audio for surround sound.

- System Layer

  Its system layer is defined in ISO 13818-1 and ITU H.222.0.

  It provides audio and video synchronization through time stamps and can carry program and transport streams.

1.6.5 MPEG-3

MPEG-3 was planned to support HDTV applications with dimensions of up to 1920x1080x30 Hz. However it was discovered that the MPEG-2 standard worked very well for HDTV rate video.

Now HDTV is a part of the high-1440 level and high level defined for MPEG-2

1.6.6 MPEG-4

MPEG-4 is the current MPEG project being developed to provide enabling technology for many applications that are becoming possible through the Internet, such as interactivity, physical network independence and decoding downloadability.

It has reached working draft level in 1996 and should reach international draft level in the end of 1998.

MPEG-4 is oriented towards very low bandwidth applications (around 64 kbps), such as videophones, and video resolution of 176x144x10 Hz.

1.6.7 Overview of ITU Standards

The H-series of the ITU-T (formerly CCITT) standards focus on audiovisual communication services. This family of standards is subdivided into video, audio and data specifications. The video standards include H.261, H.262 and H.263. Audio is specified in G.711, G.722, G.723 and G.728. Data standards are defined in T.120. These recommendations are very detailed and lengthy, so we do not cover them here.

Here is a brief description of the current related ITU standards:

- **H.320** (Narrow-band visual telephone systems and terminal equipment) is also known as px64 or ISDN Conferencing.

  In late 1990, the ITU-T approved the H.320 family of standards, for use in nx64 kbps ISDN lines. Its maximum transmission rate is 1.9 Mbps, but typically 128 kbps is used. It uses H.261 video standards, which define the use of two video formats: CIF and Q-CIF. As explained in 1.3, “Digital Audio and Video Formats” on page 5, the CIF display resolution is 352 pixels per line x 288 lines and is normally displayed full-screen. Q-CIF, with 176x144 resolution looks best when squeezed to quarter-screen. H.261 requires codecs to operate with Q-CIF, while support to CIF is optional in the standard. Under the standard, an H.261-compliant codec can code fewer than 30 frames per second, so the picture quality of different manufacturers will vary based on encoder design factors.

  H.320 is usually low-quality video and telephone-quality audio. Its system layer is defined in H.221 to carry synchronized video and audio, but it is not compatible with MPEG-2 H.222.1 transport stream. H.320 uses H.242 control protocol for the establishment of audiovisual channels between terminals. This is currently used in many desktop and room conferencing systems, including the FVC offering from IBM.

- **H.323** covers conferencing over POTS (Plain Old Telephone Service), specifically for local area network environments.

  On the LAN, video is just one of many kinds of information using the transmission resources. Therefore, desktop videoconferencing requires special compression technologies to adjust to constantly changing available bandwidth. H.320 falls short here, because the standard′s architecture is ISDN-oriented. Throughout 1996, the ITU-T was working on the H.323 to provide a gateway for LAN/WAN videoconferencing, as well as Internet videoconferencing.

- **H.324** (Terminal for low bit rate multimedia communication) is geared towards PC-based video phones and stand-alone video phones.

  H.324 calls for transmission of real-time video, audio and data over V.34 modems, which have a maximum bit rate of 22.8 kbps to 33.6 kbps. Part of the H.324 is a newer video standard called H.263, which defines and requires a new SQCIF picture format of 128 x 96 pixels. QCIF is also required and CIF support is optional. H.324 is important in delivering desktop videoconferencing products to the home and small office (SOHO) market.

- **H.310** (Broadband communication systems and terminals) is a superset of H.320.

  Its video portion is defined in H.261 and H.261. MPEG-2 streams of Simple Profile @ Main Level and Main Profile @ Main Level are supported.

  The audio options include MPEG-1 and MPEG-2 compression as well as G.7xx audio compression.

  The system layer is defined in H.222.1 and ISO 13818.1. It uses MPEG-2 transport stream and defines specific code points for ITU usage.

  The control protocol is defined in H.245 for establishment of audiovisual channels between terminals. It determines compatibility between terminals and uses the best operational features of terminals through the exchange of control information.

1.6.8 Conclusion

The graph in Figure 11 on page 23 shows the behavior of several compression methods, measuring their video quality levels against a range of transmission rates.

The dashed line on the graph indicates a boundary of the quality levels. Above this line, the CCIR-601 video format is used for high-quality video. Below the line, the SIF format is used and medium-quality video is achieved.

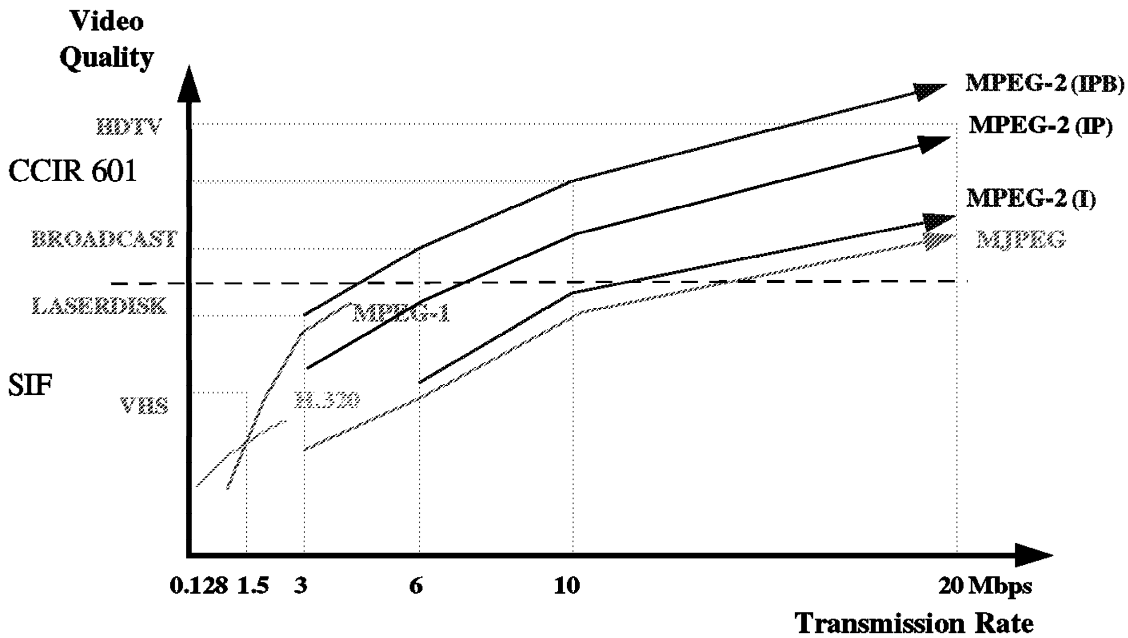

Figure 11. Transmission Rate vs Video Quality

Different levels of quality are marked, from VHS to HDTV. For instance, to achieve broadcast-quality, the graph indicates that MPEG-1 cannot be used. MPEG-2 would require a transmission rate from 6 Mbps (with I, P and B-picture compression) to almost 20 Mbps (with only I-picture compression), and M-JPEG would need more than 20 Mbps.

Higher transmission rates provide higher quality levels when the same compression technique is used.

When comparing different methods, we should first look at the quality level and then examine the bit rate required to achieve that quality. Usually different bit rates are required for different compression methods to achieve the same video quality.

Chapter 2. ATM Overview

This chapter provides an overview of ATM principles and terminology which may be of use to you in understanding the concepts presented in this book.

2.1 ATM and the B-ISDN Protocol Reference Model

Broadband Integrated Services Digital Network (B-ISDN) is the technology that is being developed to give users a single service to support voice, video and data.

ATM was chosen as the transfer mode technology to deliver B-ISDN. Consequently, you may see the term B-ISDN and ATM used interchangeably.

Part of the initial ITU-T recommendation on B-ISDN included the B-ISDN Protocol Reference Model, which specifies a layered architecture that defines basic principles and characteristics of B-ISDN. The model also defines a series of planes. These are referred to in ATM standards, for example those produced by the ATM Forum.

The *user plane* provides for the transfer of user application information. It contains a physical layer, an ATM layer and multiple ATM adaptation layer required for different service users (for example, constant bit rate (CBR) and variable bit rate (VBR) services).

The *control plane* deals with connection setup and release and other connection control functions necessary for providing switched services. The control plane structure shares the physical and ATM layers with the user plane, and also includes ATM adaptation layer (AAL) procedures and higher layer signaling protocols.

The *management plane* provides specific functions and also the capability to exchange information between the user plane and the control plane. It contains two sections: layer management and plane management. Layer management performs layer-specific management functions while the plane management performs management and coordination functions related to th complete system.

Asynchronous Transfer Mode (ATM) is the new transmission technology developed to be used from the desktop-to-desktop across worldwide information networks. ATM technology was developed to handle different types of information, including voice, video and data traffic. It is radically and fundamentally different from previous technologies, using cell switching instead of packet switching or shared media solutions.

One of the driving forces behind ATM was the need to exploit evolving telecommunications capabilities. These capabilities are characterized by the following:

- High-speed communication links (up to 10 Gbps) are available using fiber technology.
- Extremely low error rate. Compared to traditional copper communications facilities, optical fiber is much better.

- The bandwidth of fiber cable is not limited in comparison to the networking equipment using it. The installation of fiber cables is, therefore, a very economical, long-term investment.

At the same time high bandwidth requirements have been evolving in existing LAN networks. This evolution is associated principally with the increasing day-to-day use of client/server technology, which has lead to the following commonly encountered problems for network managers:

- In distributed client/server environments servers may require bandwidth beyond the limit of current LAN technologies.
- Applications require more and more bandwidth, so to maintain Quality of Service, fewer and fewer workstations can be connected to the same shared LAN segment.
- High segmentation of a large network into segments and subnets solves the problems described above, but increases the complexity of the network, increases the latency and delay of the overall network, and increases the costs of the network in installation, management and maintenance time.
- Multimedia applications require isochronous communication that is not adequately handled by shared media LANs.

Developing ATM switching technologies has also meant that existing communication practices and protocols are no longer efficient:

- The network and individual switches must be capable of switching data at the full link speed, so existing software switching methods are no longer possible.
- There is no longer time for error recovery at the physical layer, and low error rates on links make it impractical at the network level; therefore it is more effective to check data integrity at the application level, and retransmit the whole message if required.
- Existing flow and congestion control mechanisms require very significant processing time, which is no longer available in very fast networks. Flow and congestion controls for data within conventional networks must be replaced with control as data enters the network.

ATM has addressed the application, telecommunication and protocol issues highlighted above through the development of a series of standards currently being implemented by vendors. In summary the following are some of the benefits of this ATM technology:

- ATM handles different types of traffic (voice, video, data, image, multimedia, etc.) in an integrated way.
- ATM can be used in both LAN and WAN environments, providing a basis for end-to-end internetworking across enterprises.
- ATM uses new hardware switching technology, which allows very fast campus networks to be built.
- ATM is a very cost-effective technology for building a campus network, because users can connect to the network using adapters that support bandwidth according to their individual requirements. Workstations can be connected with low-speed adapters, while servers and backbone switches can use high-speed connections.

- ATM is open. It is defined by a consortium of vendors and users (the ATM Forum) and standardized by the International Telecommunication Union - Telecommunication (ITU-T, formerly CCITT).

2.2 Basic ATM Terminology

The following sections provide you with a brief introduction to some basic ATM terminology.

2.2.1 ATM Network Characteristics

From a design point of view it is important to understand the fundamental characteristics of an ATM network, which are as follows:

2.2.1.1 Connection-Oriented

An ATM network uses connection-oriented technologies. This means that there is no way to send data through the ATM network before a connection is established. A connection may be either *switched*, in which case it is built call-by-call by the signaling procedure, or *permanent*, that is, pre-established in the network based on pre-configured information.

2.2.1.2 Connectionless Operation

The majority of pre-ATM networks operate in connectionless mode. Several methods have been defined to emulate connectionless data transfer through connection-oriented ATM networks for compatibility with existing applications. These solutions operate above the ATM transport layer and rely either on the ATM connection being established by the first cell of the data or on using pre-defined permanent ATM connections.

2.2.1.3 Guaranteed In-Sequence Delivery

Because ATM cells are delivered over a virtual connection, each cell travels along the same route, and all cells will be transferred to the destination end station in the same order as they were presented to the network. This restricts the network to using a single path for cells on a single virtual connection, even though other existing physical channels may be available or underloaded.

2.2.1.4 Broadcast and Multicast

Although ATM is connection-oriented, point-to-multipoint connection types (multicast) are defined. These can be used to send data simultaneously to more than one end system and use a tree structure as illustrated in Figure 12 on page 28.

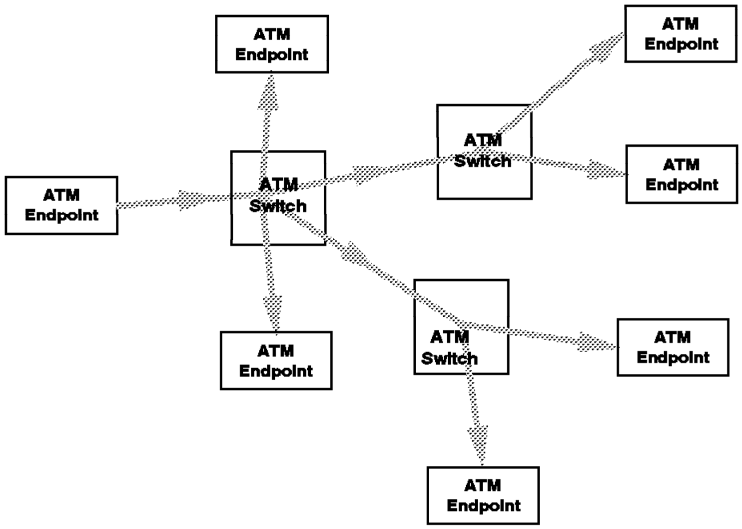

Figure 12. Multicast Tree

Point-to-multipoint connections are first established as a single point-to-point connection between the root end system and one leaf. Once this connection is established, a second leaf is connected to it using the optimal route from the established connection. This algorithm is used until the last leaf is connected to the tree.

The following are some characteristics of a multicast connection:

- Communication is available from the root to the leaf.
- Data may be sent from leaf to the root, but does not allow leaf-to-leaf communication over this connection.
- The multicast tree may be set up by signaling or by the network administrator as a permanent connection.

2.2.1.5 Quality of Service (QoS)

Each ATM virtual connection has Quality of Service characteristics associated with it. During congestion, when a network cannot recover from an overload, it discards only the cells marked as low priority. The network can select which cells to discard depending on the QoS characteristics of the virtual connection.

The QoS parameters defined by the ITU-T are as follows:

- Cell Transfer Delay (Network Latency)
- Cell Delay Variation (Jitter)
- Cell Transfer Capacity (Speed - allowed average and peak rates)
- Cell Error Ratio
- Cell Loss Ratio
- Cell Misinsertion Rate

Each virtual path also has a QoS associated with it. The QoS of a virtual connection within a virtual path may be lower than that of the virtual path, but cannot be higher.

2.2.1.6 Cell Loss and Cell Discard

Cells may be lost or discarded by an ATM network. The network does not detect the loss of cells, and does not signal the user when it has discarded cells from on a particular connection.

Some variable bit rate applications for voice and video can produce two types of cells. Standard cells contain basic information, and optional cells contain information on an alternate quality of service. If the user equipment can mark the optional cells, it can avoid the loss of essential information during network congestion.

In fast cell based networks, congestion is handled by discarding cells, and recovery is accomplished by retransmission of the full block rather than individual cells.

2.2.1.7 Congestion Control

ATM networks do not have flow control of the kind found in traditional packet switching networks. This is because traditional windowed link protocols are no longer effective at high-link speeds.

In ATM the parameters for a connection (for example, requested bandwidth, QoS) are examined before the connection is established, and the connection only allowed if the network can support the desired parameters.

The network allocates resources on a statistical basis. It allows for the possibility that demand may exceed available network resources, in which case the network will discard cells.

2.2.1.8 Input Rate Policing

At the entry point of the network, the ATM switch monitors the rate of data arrival for a virtual connection or virtual path according to the negotiated QoS parameters for the connection. It will take action to prevent an ATM endpoint from exceeding its allowed limits using a technique called back-pressure.

In the case of overload, depending upon the network configuration, the network may either:

- Discard cells received over the allowed maximum rate.
- Mark the overloaded cells with CLP.

2.2.1.9 End-to-End Data Integrity

No end-to-end data integrity is provided by the ATM transport layer. This function is the responsibility of end user equipment or of a higher layer protocol.

2.2.1.10 Priorities

There is no priority defined within an ATM network. CLP is not a real priority mechanism, despite its name, as it simply defines which cells may be discarded in the event of congestion.

2.2.2 The Structure of ATM Network

The conceptual structure of an ATM network is shown in Figure 13. The main components of the network are described in the following sections.

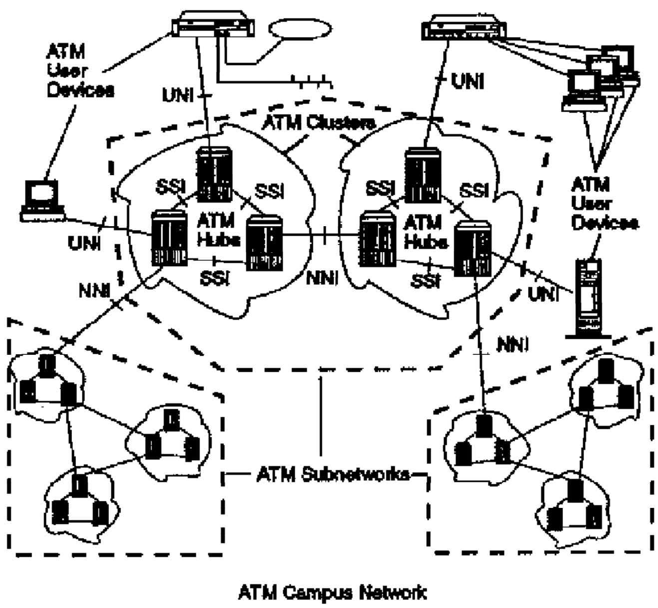

Figure 13. ATM Network Structure Overview

2.2.2.1 ATM Network

An ATM network is a set of interconnected ATM subnetworks controlled as one administrative domain. Networks can be either private and public. Individual networks can be interconnected to each other.

2.2.2.2 ATM Subnetwork

An ATM subnetwork comprises a single ATM cluster or a set of ATM clusters that are operating as a single routing domain. Clusters within the subnetwork are interconnected by NNI interfaces and share the same routing domain number (RDN) in their ATM addresses.

2.2.2.3 ATM Cluster

An ATM cluster comprises a single ATM subsystem or a set of ATM subsystems interconnected by SSI interfaces. Each subsystem in an ATM cluster is configured to have the same ATM cluster number (ACN) in its ATM address.

2.2.2.4 ATM Subsystem

An ATM hub is an example of an ATM subsystem, which typically carries out the switching function in a network. Each subsystem in an ATM cluster is identified by a unique hub number (HN) in its ATM address.

2.2.2.5 ATM User Device

An ATM user device is the end system that encapsulates data into ATM cells and forwards them to a local switch for forwarding over the network. Each user device has a unique ATM address comprising a network prefix, routing domain, cluster number, hub number in addition to a local n byte suffix. Typically, the user device learns the remainder of its address (all but the local suffix) during initialization using facilities of the user device to switch interface, called the UNI interface.

2.2.3 ATM Connections

ATM differs from existing LAN networks, because it uses connection-oriented technology. The connection, in ATM terminology, is a point-to-point or point-to-multipoint link from one end system to another across a series of ATM switches in a network.

This connection-oriented technology simplifies the routing of cells across the ATM network. Station destination and source addresses do not need to be carried in each ATM cell; only the connection identifier is required by each ATM switch to route the cell correctly. Because the information between end systems is sent over the single route described by the connection identifier, the information is received in the same order as it was sent. This in-sequence delivery is required especially for voice and video traffic and also simplifies the processing of data traffic.

Before data can be transferred, a virtual connection needs to be established between the end systems either by using a pre-determined fixed path or by means of the signaling protocol. The connection may, therefore, be permanently established by the network operator (*Permanent Virtual Connection, or PVC*), or temporarily established on request from an ATM end system by signaling (*Switched Virtual Connection, or SVC*).

To handle each connection, ATM uses the concept of virtual paths and virtual channels. These are described in the following sections:

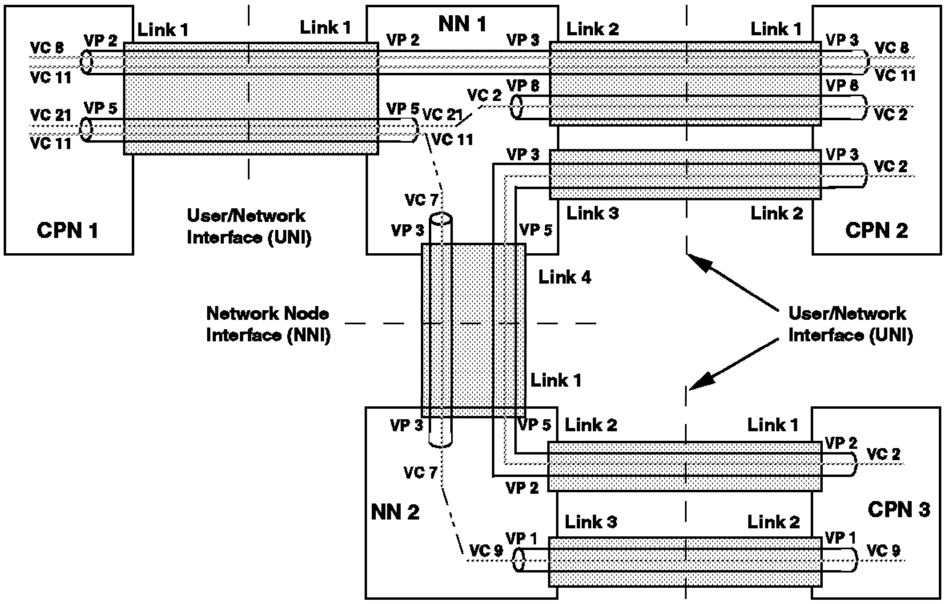

Figure 14. Routing Concept in an ATM Network

2.2.3.1 Virtual Path (VP) and Virtual Path Indicator (VPI)

As illustrated in Figure 14 a virtual path is an aggregate route through a network representing a group of virtual channels (VCs). VPs may exist:

- Between ATM end systems
- Between ATM end systems and ATM switches
- Between ATM switches

The virtual path indicator is the indication of the virtual path to be used by a cell and is contained within each cell in the network.

2.2.3.2 Virtual Path Link (VPL)

A virtual path link exists between the points where a VPI value is assigned or where is translated or determined. Typically these points would be switches in the ATM network.

2.2.3.3 Virtual Path Connection (VPC)

A virtual path connection is the concatenation (sequence) of VPLs that extends between virtual path terminations.

2.2.3.4 Virtual Path Switch (VPS)

A virtual path switch is the processing function that connects VPLs to form VPCs. This function translates VPI values (label swapping) and directs cells to the correct output link at a particular ATM switch.

2.2.3.5 Virtual Path Terminator (VPT)

The virtual path terminator is a a processing function also. It terminates each VP and makes the associated VCs available for separate and independent connection routing.

2.2.3.6 Virtual Path Connection Identifier (VPCI)

This identifier of a VP connection is returned by the ATM network when call setup is performed by a user device. It is 16 bits long, and used by the signaling protocol instead of the VPI, which is unique only within a single ATM link.

2.2.3.7 Virtual Channel (VC) and Virtual Channel Indicator (VCI)

A virtual channel is defined in ATM as an unidirectional connection between user devices.

The virtual channel indicator (VCI) is the indication of the virtual channel to be used by a cell and is contained within each cell in the network.

2.2.3.8 Virtual Channel Connection (VCC)

A virtual channel connection is the end-to-end connection along which a user device sends data. Since virtual channels are unidirectional, a VCC would normally consist of two virtual channels to provide full full-duplex data transfer.

2.2.3.9 Virtual Channel Link (VCL)

A virtual channel link exists between the points where a VCI value is assigned or where it is translated or determined. Typically these points would be switches in the ATM network.

A virtual channel link is a separately defined data flow within a link or virtual path. A virtual channel connection (VCC) through a network is a sequence of interconnected (concatenated) VCLs.

The relationship between links, VPs and VCs is summarized in Figure 15 on page 34 and Figure 16 on page 34.

2.2.3.10 Virtual Channel Switch (VCS)

The virtual channel switch is the VC switching function shown in Figure 16 on page 34, where VCLs are connected together to form VCCs. To do this they terminate VPCs and translate VCI values.

There are limits on the number of VPs, VCs, etc. within an ATM network as follows:

- The maximum number of VPs on links is determined by the number of bits allocated to address the VPs in the cell header (VPI). This is either 8 or 12 bits.
- The maximum number of VCs within a VP is determined by the number of bits allocated to address the VCs in the cell header (VCI). This is 16 bits.
- There may be additional limits imposed by the capacity of specific ATM equipment.

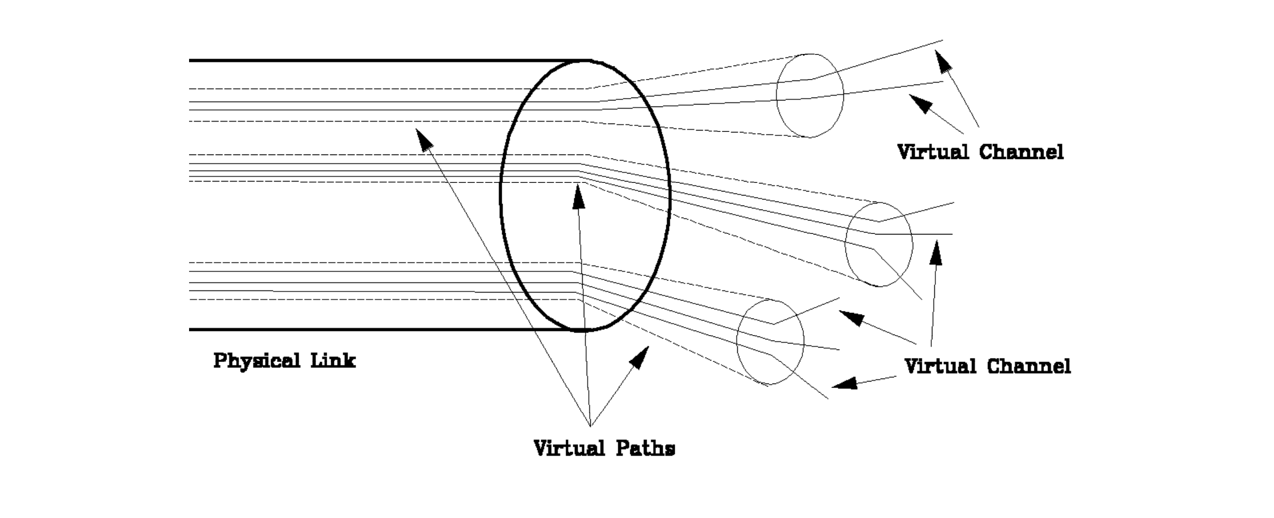

Figure 15. Link, Virtual Path and Virtual Channel Relationship

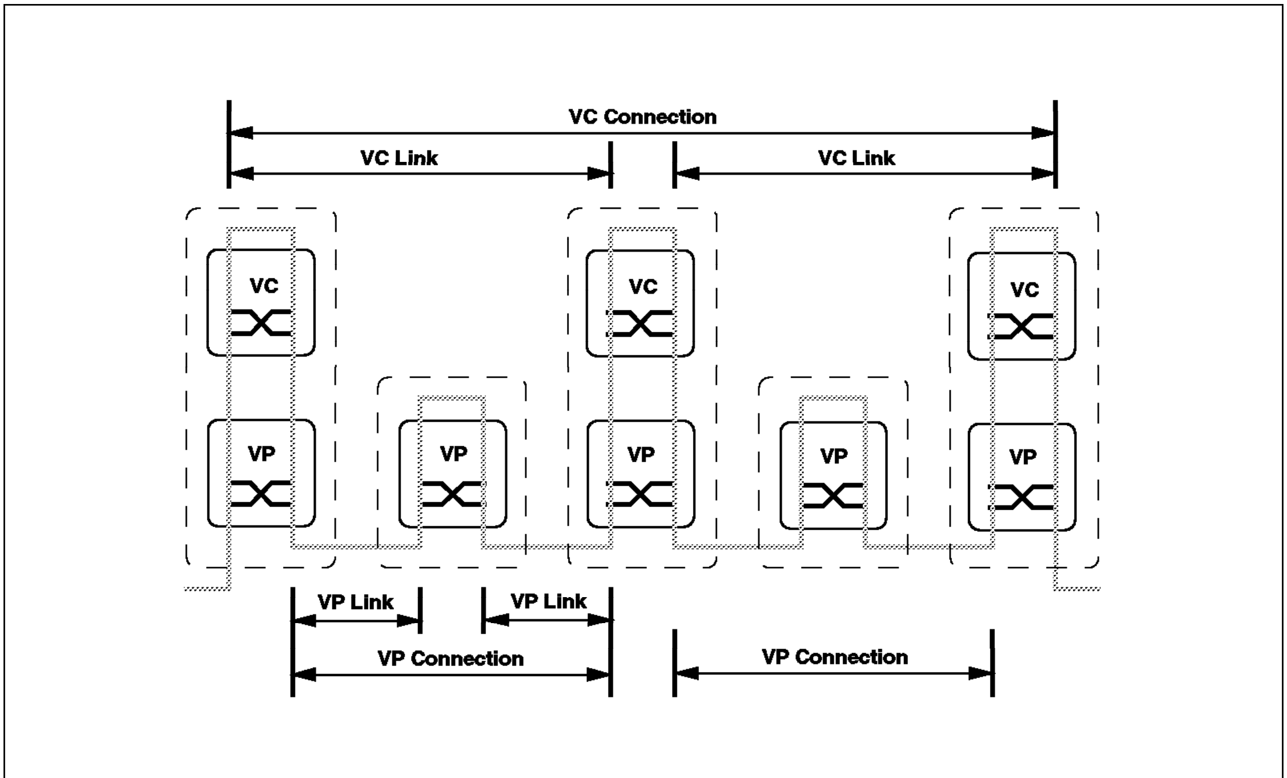

Figure 16. VP and VC Concepts

2.2.4 Routing/Switching ATM Cells

An ATM cell is transmitted along a virtual channel connection according to the routing information contained in its header. This routing information is swapped at every switch along the path of the connection, to enable the routing of the cell to the next switch along the connection. This process is referred to as label swapping.

The routing information in the cell consists of the VPI and the VCI fields in the cell header. The definition of the VPI/VCI mapping is established when the connection is established. The mapping for the connection is held at every intermediate switch and consists of VPI/VCI input fields mapped to VPI/VCI output and port output fields.

This means that each VPI/VCI pair is associated with a particular port on a switch and each VPI/VCI associates a cell with an input link and a corresponding output link. Figure 17 shows an example of mapping VPI/VCI. Based on the VPI/VCI in the header of the cell an ATM switch can identify the output link across which the cell is to be routed and give the new link identifiers to the cell.

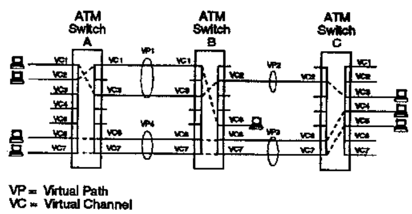

Figure 17. VPI and VCI Switching in an ATM Switch

Switching of all virtual channels within a virtual path can take place inside a switch. In this case the switch unpacks the virtual path into unique virtual channels, and using its routing table groups them together again, forming new outgoing virtual paths as shown in the top portion of Figure 18.

The bottom portion of Figure 18 shows an example where the VPI is switched without unpacking it.

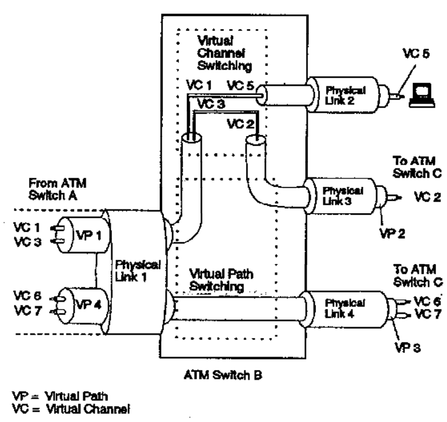

Figure 18. VPI and VCI Mapping Across an ATM Switch

2.2.5 ATM Cells and Cell Format

In ATM networks, information is transmitted in cells. Cells are fixed-length packets. Each packet is 53 bytes long; 48 bytes is the *payload*, with a 5-byte *header*. The header contains control information, including the VPI/VCI route identifier that defines cell route information. The header is error-checked to avoid errors being propagated over the network. The payload is user information which is not protected by error checking at the ATM network level.

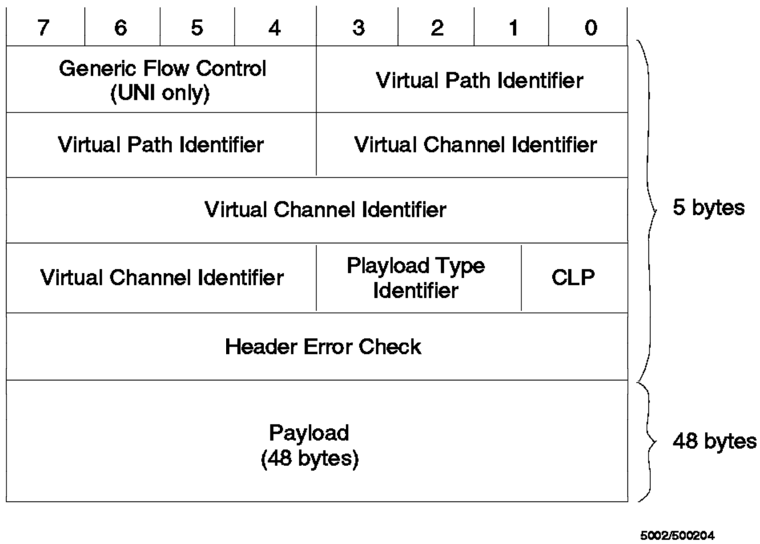

Figure 19. ATM Cell

The cell formats for the user network interface (UNI) and the network node interface (NNI) are illustrated in Figure 20 and Figure 21 on page 37. The fields within the cell are described below:

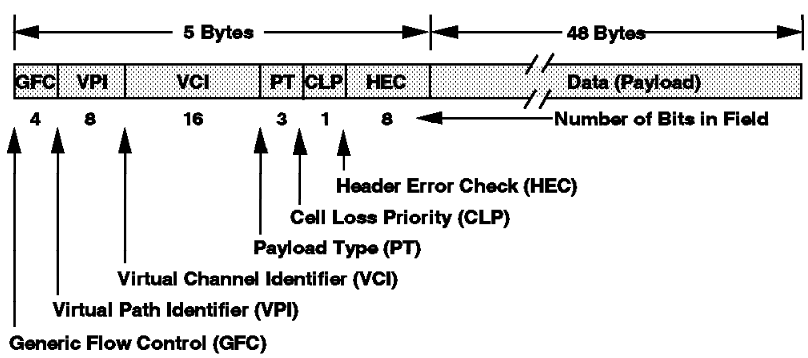

Figure 20. ATM Cell Format at the User Network Interface (UNI)

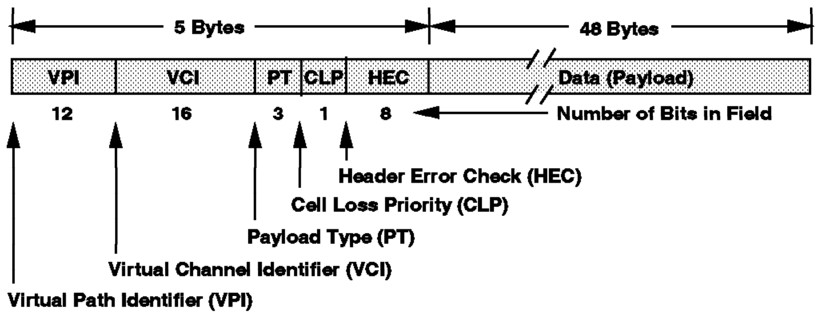

Figure 21. ATM Cell Format at the Network Node Interface (NNI)

2.2.5.1 Cell Size

An ATM cell always contains 48 bytes of data, with a 5 byte header. The cell size was defined by the ITU-T as a compromise between data and voice requirements.

2.2.5.2 Generic Flow Control (GFC)

Currently use of the GFC has not yet been standardized, but it is defined as having the same length as in the UNI cell.

Note: This field is missing in the NNI cell header so the flow control can only be a procedure local to the user device.

2.2.5.3 VPI and VCI

These two fields, Virtual Path Identifier (VPI) and Virtual Connection Identifier (VCI), are the most important fields in the cell header. Together they identify the logical connection (the virtual connection) over which the cell is travelling. Some VPI/VCI values are reserved for signaling (for example connection establishment), and for maintenance and resource management.

If the VPI and VCI values are set to zero, it means that the cell is empty. Empty cells may be required in a network to maintain physical link protocols.

2.2.5.4 Payload Type (PT)

This is a 3-bit field. Bit 0 determines if the cell is user data (Bit 0 = 0) or operations, administration and maintenance (Bit 0 = 1).

Bit 1 is used when the cell carries user data. Bit 1 = 1 means that congestion was experienced somewhere along the route passed by the cell.

Bit 2 is used by higher layer processing. Bit 2 = 1 means that the cell is the last cell of the user data frame. If Bit 2 = 0, the the cell is the first or the middle cell of the frame.

2.2.5.5 Cell Loss Priority (CLP)

When set, this bit indicates that the cell is a low priority cell. When the system needs to discard cells in a congested situation, these cells should be discarded first.

2.2.5.6 Header Error Check (HEC)

This field allows the correction of all single-bit errors or the detection of multi-bit errors in the header part of the cell.

2.2.6 ATM Signaling

ATM signaling is the process used for dynamic setup and clearing of ATM switched virtual connections (VPCs and VCCs) at the UNI interface.

Permanent virtual connections are established by a network operator. If a connection is permanent, and if the network fails and is restarted, the circuit is reestablished. In case of recovery from failure the network will not re-establish lost switched virtual connections.

The key elements of ATM signaling are the following:

- Signaling takes place on separate VCCs from those used by user data. The same principle is used as in narrowband ISDN where the D channel is used to initiate connections.
- There is a method to set up additional signaling channels besides the predefined channels (although this is not used in current ATM implementations).
- Point-to-point signaling is the default signaling method, but broadcast signaling may be implemented in the future. Point-to-multipoint connections are established using point-to-point signaling.
- Class B (AAL-2) services have not yet been defined, so it is not supported by the signaling protocol.
- The Class D protocol is also not supported, because clients must have connections to connectionless servers, so the calling procedure is not required.
- Connection setup by third party equipment, who is not involved in the data transfer is not supported.

The routing framework developed for the routing of ATM connections is based on an extension to the OSPF (Open Shortest Path First) routing protocol known as Widest Path OSPF. In this modification, the cost of a link is defined as its available bandwidth. This algorithm will select the path providing the widest bandwidth for the connection, instead of the shortest as defined in “traditional” OSPF.

These characteristics allow the connection path to be computed in advance (pre-computation), hence the route computation process is independent of the connection setup process. During connection setup, the only checking required is whether there is available bandwidth for the new connection according to its QoS characteristics. Using precomputed routes connection setup time will be shorter and a higher connection setup rate can be provided by the ATM network.

The following functions and subfunctions are defined for ATM signaling:

1. Call Establishment

- Setup
- Call Processing
- Connect
- Connect Acknowledge

2. Call Clearing
   - Disconnect
   - Release
   - Release Complete
3. Status
   - Status Enquiry
   - Status
4. Point-to-Multipoint Messages
   - Add Party
   - Add Party Acknowledge
   - Add Party Reject
   - Drop Party
   - Drop Party Acknowledge

Typical connect and disconnect procedures are shown in Figure 22 and Figure 23 on page 40.

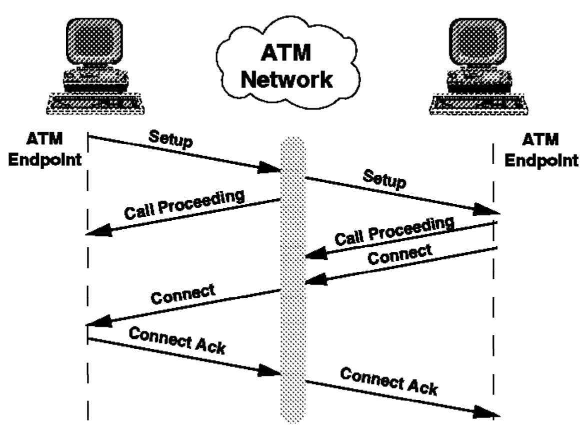

Figure 22. ATM Switched Circuit Setup Protocol

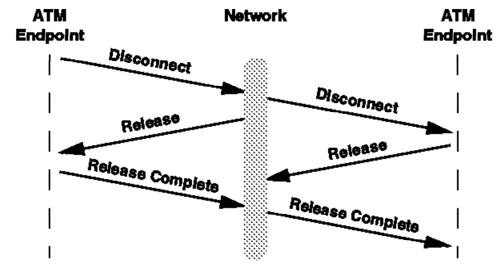

Figure 23. ATM Switched Circuit Disconnect Protocol

2.2.7 ATM Address Format

The address formats used in ATM are shown in the Figure 24.

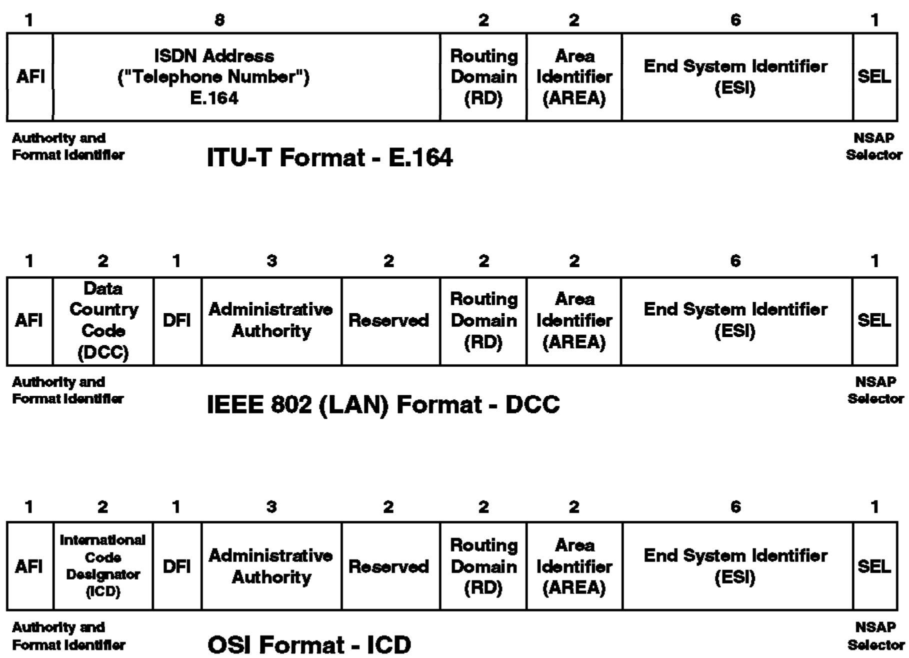

Figure 24. ATM Address Formats

The meaning of the individual fields in the address are shown in the figure.

The address must be separated into the following two distinct parts:

- The end system address identifies the end system uniquely in a cluster. The field is 7 bytes long and consists of an ATM end system address (similar to a

LAN MAC address) and a single byte end system identifier (ESI) that identifies a subcomponent of an end system.

- The remaining 13 bytes contain the network part of the address that includes fields for hub number, cluster number, routing domain as well as the standard network prefix.

The following are the three different formats for ATM addresses, each controlled by a different authority:

1. ITU-T (E.164) Format: This format is essentially the same as telephone style addressing. It is specified by the ITU-T and will be used by public (carrier provided) ATM networks.
2. DDC (Data Country Code) Format: This format carries LAN addresses as specified in IEEE 802 recommendations.
3. IDC Format: This format is specified by the ISO for OSI.

The ATM Forum specifies that equipment in a private network must support all three formats.

2.2.8 ATM Adaptation Layers (AAL-s)

The network characteristics required by various types of traffic over an ATM network are provided by an ATM adaptation layer which is found in each end system and, in special forms, in switches also.

The ITU-T has defined four different generic service classes of network traffic, each of which must be treated differently by an ATM network. These classes are designated Class A to Class D and four different types of ATM adaptation layers (AAL) have been defined to realize the necessary network characteristics to handle them. The relationship between these is shown in Figure 25 on page 42, and the functions of the appropriate layers is in Figure 26 on page 42.

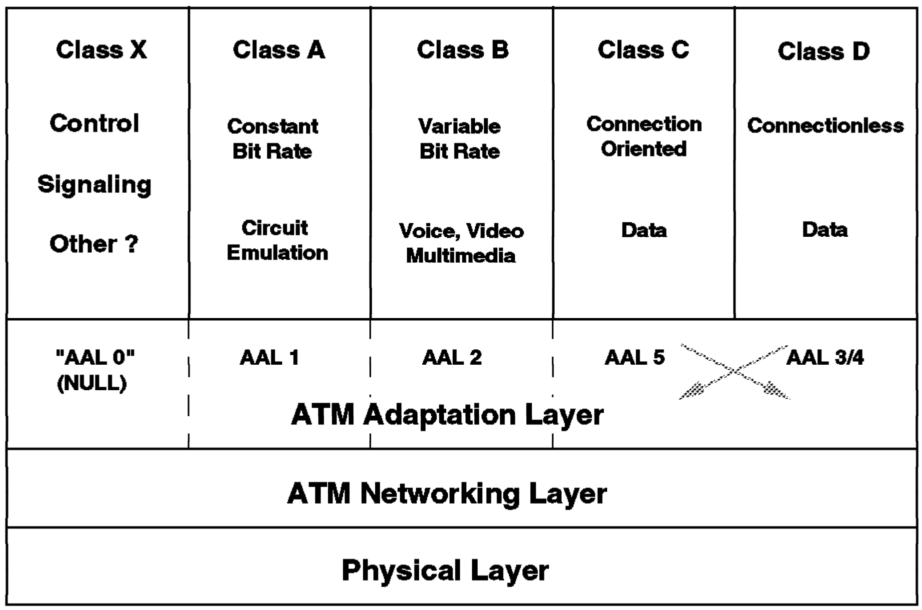

Figure 25. Service Classes and AAL Types

Layer Name		*Functions*	
Higher Layers		Higher Layers Functions	
AAL	Convergence sublayer	Service specific (SSCS) Common part (CPCS)	Layer Management
	SAR sublayer	Segmentation and reassembly	
ATM		Generic flow control Cell header generation/extraction Cell VPI/VCI translation Cell multiplexing/demultiplexing	
Physical	Transmission convergence (TC) sublayer	Cell rate decoupling Cell delineation Transmission frame generation/recovery	
	Physical medium dependent (PMD) sublayer	Bit timing Physical medium	

5002/500205

Figure 26. Composition of the ATM Sublayers

The following sections provide a short description of these service classes:

2.2.8.1 Class A (Circuit Emulation)

This service emulates a leased line and is used for traffic that has a constant bit rate, for example voice and video.

The characteristics of Class A are the following:

- A constant bit rate at source and destination
- A timing relationship between source and destination
- A connection between end users of the service

To realize these functions, the adaptation layer must perform the following services:

- Segmentation and reassembly of data frames into cells
- Handling (by buffering) of cell delay variations
- Detection and handling of lost, discarded, misrouted or duplicated cells
- Recovery of the source clock frequency
- Detection of bit errors in the user information field

2.2.8.2 Class B (Variable Bit Rate Services)

This service is intended for isochronous voice and video traffic which may be coded as variable rate information, and requires a timing relationship between the ends of the connection. The service is strictly connection-oriented.

The services provided by the Class B adaption layer are the following:

- Transfer of variable rate information between endpoints.
- Transfer of timing between source and destination.
- No indication is provided of lost or corrupted information.

Some video applications (such as videoconferencing) require synchronization between voice and video, while others can be transmitted in multicast mode (just like a film) and are not sensitive to network delay.

To design and control a network for Class B traffic is very challenging because of the unpredictable way high bandwidth is required, and because the data rates are often near to the peak rate capability of the network.

2.2.8.3 Class C (Connection-Oriented Data)

Class C traffic is traditional data traffic, such as SNA and X.25. The service offered for it is connection-oriented and it supports variable rate information flow.

The services provided by the Class C adaptation layer are the following:

- Segmentation and reassembly of frames into cells
- Detection and signaling of errors in user data frames
- Possible multiplexing and demultiplexing of multiple end user connections into a single ATM connection

2.2.8.4 Class D (Connectionless Data)

The Class D service is connectionless and also supports variable rate information flow. It is intended to support connectionless protocols such as TCP/IP.

The services provided by the class D adaptation layer are the following:

- Segmentation and reassembly of frames into cells
- Detection and signaling of errors in user data frames
- Multiplexing and demultiplexing of multiple end user connections into a single ATM connection
- Network layer addressing and routing

2.2.8.5 Class X (User Defined)

This is a connection-oriented ATM transport service where the network characteristics are user defined. Only the required bandwidth and the QoS parameters are used by the network.

As shown in the Figure 25 on page 42, the different service classes are represented by the appropriate ATM adaptation layer types.

2.2.8.6 AAL-0

This is the null AAL and corresponds to a process that connects the AAL service interface directly to the ATM networking service.

2.2.8.7 AAL-1

AAL-1 is used for Class A (constant bit rate) traffic. In practice this may take the form of data in an SDH or PDH frame where the frame rate is constant but data exists in a specific part of the frame, so that it arrives at the network in short periodic bursts.

When these cells are transported through the network, dependent upon the current load on the network, they may suffer relative delay. This delay is not consistent, and but can introduce jitter. The receiver has to buffer the received data to avoid problems of overrun or underrun of data frames, with consequent additional delay in the data stream.

Figure 27 shows the format of AAL-1 cell.

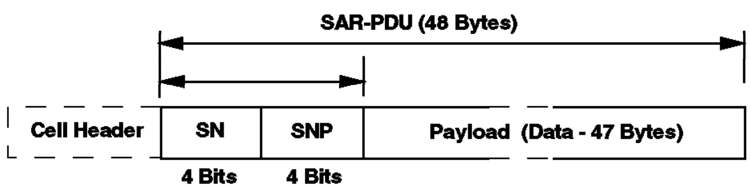

Figure 27. Cell Format for AAL-1

There are the following two additional control fields in the cell, but because there is no additional space in the cell header, they are located at the beginning of the payload. Each field is 4 bits long.

Sequence Number (SN): This field contains a 3-bit sequence number field that operates in cyclic fashion, with a 1-bit CS indicator (CSI bit), which is used depending upon the type of the traffic.

Sequence Number Protection (SNP): This field is CRC protection for the SN field.

In order to minimize the delay caused by assembly and playout, it is possible to send cells that are not full. This fact must be specified between the end systems at the time of circuit establishment.

2.2.8.8 AAL-2

AAL-2 is used for Class B traffic. AAL-2 processes data streams such as AAL-1, the only difference is the variable bit rate due to compression. One of the key problems of AAL-2 is how to handle the skew between voice and video information.

AAL-2 is currently absent from the draft standards. The detailed description is being delayed because of the above mentioned problems.

2.2.8.9 AAL-3/4

AAL-3/4 is used for Class C and Class D traffic.

It is a relatively complex, high-function AAL that will offer assured data delivery.

When the AAL detects corruption of a data frame, it will be able to re-transmit the data frame, although the details of this operation have not been defined.

In blocking mode short data frames are blocked into a longer data unit for transmission through the network, and the multiplexing of several AAL-to-AAL connections onto a single VCC connection is possible.

Point-to-multipoint connections are also possible using AAL-3/4.

In this mode several additional control information bits are placed in the payload field. The format of the cell for AAL-3/4 is shown in the Figure 28 on page 46.

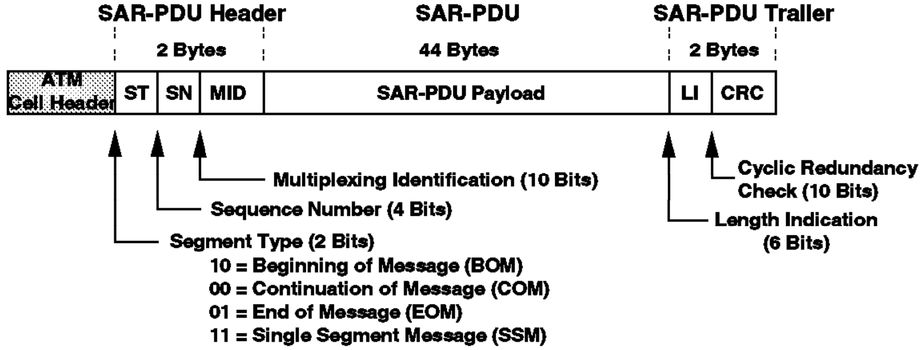

Figure 28. Cell Format for AAL-3/4

The fields have the following meaning:

Segment Type (ST): These two bits indicate where the content of this cell is located in the data frame.

Sequence Number (SN): This 4-bit field is a Module Four counter that verifies the correct sequence of the cells.

Multiplexing Identification Field (MID): In the case of multiplex mode this field indicates the connection to which the cell belongs.

Length Indicator (LI): This field specifies how much data is in the variable part of the cell. All except the last cell of the data frame should be full.

Cyclic Redundancy Check (CRC): This is a polynomial CRC to protect the whole of the cell payload except the CRC itself.

2.2.8.10 AAL-5

This AAL is often called SEAL (Simple and Efficient Adaptation Layer). It is designed to operate significantly more efficiently than AAL-3/4 but, with the exception of connection multiplexing, has the same functions as AAL-3/4.

In the case of AAL-5 the whole user frame is protected by CRC, and there is no length field in every cell as in case of AAL-3/4. In this way the whole 48 byte of the payload can be used for data transmission. AAL-5, therefore, can recognize corrupted or missing data only when the whole data frame is received at the destination end system.

2.3 ATM Network Interfaces

ATM interfaces define interoperability and connectivity between different components of an ATM network.

They have been defined by both the ITU-T and the ATM Forum. The ITU-T has been working on the definition of the public UNI and NNI. The interface specifications defined by the ATM Forum includes user-to-network interface (UNI), private network-to-network interface (P-NNI) and the data-exchange interface (DXI).

An ATM network comprises a set of end systems and a set of intermediate nodes (switches), all linked by a set of point-to-point ATM links. The different types of interface used over these links are shown in the Figure 13 on page 30 and described in the following sections.

2.3.1 ATM Physical Interfaces

The following items describe the existing ATM physical interfaces:

2.3.1.1 SONET and SDH

Synchronous Optical Network (SONET) is a U.S. standard for the internal operation of PTT optical fiber networks. It relates closely to a system called Synchronous Digital Hierarchy (SDH) which has been adopted by the ITU-T for the internal operation of carrier (PTT) optical fiber networks worldwide.

Traditionally, PTT networks have been built by using a cascade of bandwidth multiplexors at each end of the high-speed connection. This resulted in more and more stages of multiplexing to provide faster links, the internals of which were generally proprietary. For example, the U.S. used a different structure from Europe, and both the U.S. and Europe were different from Japan.

Both SONET and SDH, which was developed from it, eliminate the problems illustrated above by providing a standardized method of internal operation and management and worldwide compatibility, while enabling existing speeds to be accommodated. They permit many levels of multiplexing and demultiplexing to be achieved in a single step, and allow many different speed channels to be carried through the system. With this new scheme, it is possible to access low bandwidth channels without having to demultiplex the whole bandwidth stream.

The basic structure in SONET is a frame of 810 bytes, which is sent every 125 microseconds. This allows a single byte within a frame to be a byte in a 64 kbps digital voice channel. Since the minimum frame size is 810 bytes, this means that the minimum speed at which SONET will operate is 51.84 Mbps:

810 bytes x 8000 frames/second x 8 bits = 51.84 Mbps

The basic SONET frame is called the Synchronous Transport Signal Level 1 (STS1), as shown in Figure 29 on page 48.

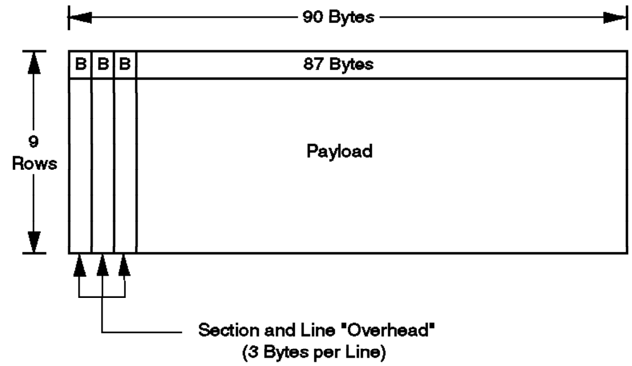

Figure 29. SONET STS1 Frame Structure

The frame is conceptualized as containing 9 rows of 90 columns each with the following attributes:

- The first three columns of every row are used for administration and control of the multiplexing system. They are called overhead in the standard but are very necessary for successful systems operation.
- The frame is transmitted row by row, from the top left of the frame to the bottom right. One frame is transmitted every 125 micro seconds.
- It is important to remember that the representation of the structure as a two dimensional frame is purely a method of showing a repeating structure. In reality it is just a string of bits with a defined repeating pattern.

Multiple STS1 frames can be bytes multiplexed together to form higher speed signals. When this is done, they are called STS2, STS3, etc., where the numeral suffix indicates the number of STS1 frames that are present. For example, STS3 is three times an STS1, or 155.52 Mbps.

For more information regarding the details of the SONET frame structure, we recommend that you refer to Appendix B3 in the *ATM Technical Overview*, SG24-4625.

2.3.1.2 SONET ″LITE″

SONET ″LITE″ is not an officially accepted term, but is a convenient description of the underlying technology upon which it is based. It is the standard proposed by the ATM Forum for use with multimode fiber in the local LAN environment, and is based on the SONET STS3c standard.

The term ″LITE″ comes from the fact that most of the management information flows of STS3c have been replaced. This is possible because of the relatively short distances involved in LAN environments which make them inherently more reliable than WAN connections. The approach of minimizing some of the frame overheads, significantly reduces the cost of implementation making it an attractive implementation option.

2.3.2 The UNI Interface

The UNI specification defines the interface between ATM end systems (such as a terminals, routers, bridges, servers, or concentrators equipped with an ATM adapter) and the ATM network. The UNI comes in two parts: the private UNI, and the public UNI.

The private UNI defines interfaces between an end system and a switch owned by an organization. An analogy would be the IEEE 802.3 specification.

The public UNI defines interfaces between an end system and a service provider′s equipment. An analogy would be the specification of the interface to public X.25.

The two interfaces have different controlling organizations: the private by the ATM Forum and the public by ITU-T.

The major differences between the two UNI specifications are as follows:

- Link Types Allowed

  Some of the link types allowed by the private UNI use protocols that work over very short distances (for example 100 meters). These would be obviously inapplicable to a public network interface.

- Addressing Format

  Public ATM networks will use the E.164 addresses (similar to those used for telephone numbers) while private networks will probably use addressing from LAN or OSI environments.

There are also some other smaller differences because of the different controlling organizations.

In the campus environment the private UNI is used, and is supported by the IBM 8260 hub. The cell format used by this interface is shown in Figure 20 on page 36.

Basic signaling information (for example, connection setup) across the UNI uses the VPI=0, VCI=5 channel. The Interim Local Management Interface (ILMI) protocol (which is part of the UNI specification) is used by switches to communicate their network address prefixes to endsystems, and for the end systems to communicate the ESI portion of their network address for the switches at initialization time. The ILMI protocol uses the VPI=0, VCI=16 channel for address registration.

2.3.2.1 Generic Flow Control (GFC) at the UNI

The header of the UNI cell contains a field named GFC. Using this field one-way flow control is defined from the ATM end system to the ATM switch. There is no control in the opposite direction.

According to the UNI standard three queues may be defined at the end system, one for uncontrolled, and two for controlled traffic. Usually, though, only one queue is used for controlled traffic. This concept is illustrated in Figure 30 on page 50.

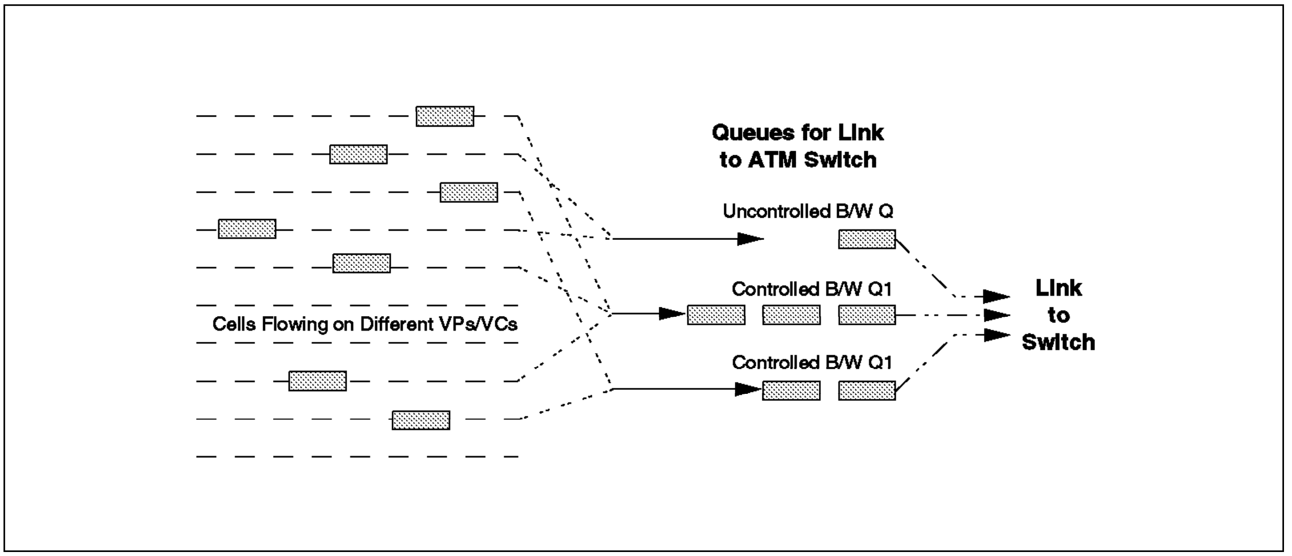

Figure 30. GFC Operation

2.3.2.2 Controlled Traffic

This is the traffic for which the GFC mechanism is defined. It is usually all the nonreserved bandwidth (NRB) traffic on the interface. Controlled traffic is distinguished in the cell header by the presence of a non-zero GFC field.

2.3.2.3 Uncontrolled Traffic

This traffic is not subject to GFC control and is treated as having a higher priority than the controlled traffic. It would normally be RB traffic for which there is a reserved bandwidth.

The flow control mechanism uses windowing. Each queue has a window that represents the number of cells that it is allowed to send before the network must respond to give permission to continue. The window is maintained as a counter in the end system. Each time a cell is sent the counter is decremented. The end system is allowed to send cells as long as the value of the counter is not zero. If the counter reaches zero the end system must stop transmitting until the counter has been reset to an initial value. During normal operation the switch sends reset signals fairly often so that the counter never reaches zero.

GFC field towards the network (outbound):

- Bit 0 is unused.
- Bit 1 indicates that the cell is flow-controlled by Q1.
- Bit 2 indicates that the cell is flow-controlled by Q2. (In the case where only one queue exists this bit is always zero.)
- Bit 3 indicates if the equipment is controlled (1) or not controlled (0).

GFC field away from the network (inbound):

- Bit 0 means HALT (1) or NOHALT (0). If HALT, the network is unable to receive input from the end system, even uncontrolled reserved bandwidth traffic.
- Bit 1 when set means reset the counter of Q1.

- Bit 2 when set means reset the counter of Q2. If Q2 does not exist, it must be 0.
- Bit 3 is reserved for future use.

2.3.3 System-to-System Interface (SSI)

SSI defines the interface operating over a link between two 8260 ATM hubs in the same cluster. It has been developed from the Interim Inter-Switch Signaling Protocol (IISP) for use on IBM 8260-based ATM networks.

The SSI uses the same signaling mechanism as the UNI. Topology information flows over the VPI=0, VCI=8 channel.

2.3.4 Network-to-Network Interface (NNI)

NNI defines the interface operating over a link between two 8260 ATM clusters, which may be within or in separate routing domains. Like the SSI, the NNI has been developed from the ATM Forum IISP specification.

Only a single NNI is permitted between two 8260 clusters. The connection can be made using a direct physical connection or using the services of an ATM provider.

Operator invention is required in order to manage network topology across NNI interfaces, because currently topology and route selection information is local to a cluster.

2.3.5 Private Network to Network Interface (PNNI)

PNNI has been derived from existing protocols and developed to allow the construction of large ATM networks. PNNI can be used either as an interface between different ATM networks or as an interface between ATM switches within a single ATM network. The PNNI specification replaces IISP which was developed purely as an interim standard.

PNNI defines all node-to-node (switch-to-switch) interfaces, hence all internal operations of an ATM network. The uses of PNNI are illustrated in the Figure 31.

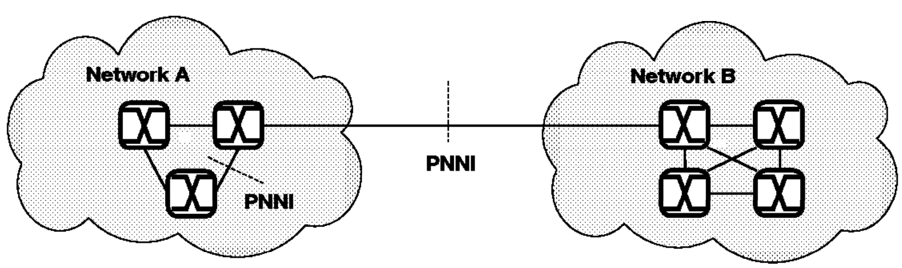

Figure 31. PNNI Usage

2.3.5.1 PNNI Functions

The basic function of the PNNI is to allow setup of VP and VC connections end-to-end through an ATM network. This means, that when an end system requests a new connection, the optimum route must be selected and the connection set up. In order to do this, the PNNI has to provide functions for the following:

1. Topology and Routing: In order to choose the optimum route through a network, information is required about the overall network topology and the current state of individual links. The protocol used to distribute topology and link state information in PNNI is called routing. In PNNI every switch carries a topology database and network directory sufficient to do basic routing.
2. Path Selection and Signaling: Because each node′s network topology detail and may be out of date, the route calculated by the first switch may not be the optimum. The route is optimized as the connection setup request is forwarded and processed by switches along the selected route.

In PNNI the switch connected to the end system making the connection request computes the best available route in advance. During call setup signaling, the originating switch sends a setup message which proceeds hop-by-hop (switch-by-switch) along the preplanned route as in the case of source routing. Each switch has the potential to update the originally selected route, if its path tables indicate that a superior one is available.

- Traffic Management: Superficially, PNNI is not concerned with traffic management, only with the management of protocols operating across it. However, at the time of connection setup the state information of links (for example, allocated bandwidth) must be taken into account. This is an example of traffic management.
- Network Management: Network management for PNNI is being developed as part of the wider suite of network management specifications and does not form part of the PNNI specification itself. However, management information bases (MIBs) that apply to PNNI are defined within the specification.

2.3.5.2 Building the Network Topology

In order to plan the route for call setup through the network the originating switch must have the following information available:

- A network directory function to enable the switch to locate each destination end system.
- A network topology map to enable the switch to calculate possible routes to the destination end system.
- Up-to-date link state information to enable the switch to determine if the desired QoS characteristics can be met on a particular route under consideration. This is normally a part of the topology map.

This information does not necessarily have to be stored within the switch itself and could be made available to it (for example, from a server). If each switch is to maintain its own database, then there must be a protocol for building this database and for keeping it up-to-date.

This could potentially create a very large amount of network traffic and as networks get larger, the amount of additional network traffic required to keep switched databases current would grow exponentially.

A mechanism is required to limit this additional network traffic while maintaining necessary network function. This is done by structuring networks into geographical areas or peer groups, such that individual switches do not need to know about the topology of the whole network, only their local peer group.

Network Hierarchy: Topology information is maintained and synchronized within each peer group. Individual members of each peer group are elected peer group leaders. Peer group leaders then communicate one with another in logical groups and exchange topology and route summarization information between themselves, hence between peer groups. In this way switches in each peer group receive information about the overall topology of the ATM network. This process is very similar to OSPF which is used in router networks.

The concept of this hierarchy is shown in the Figure 32.

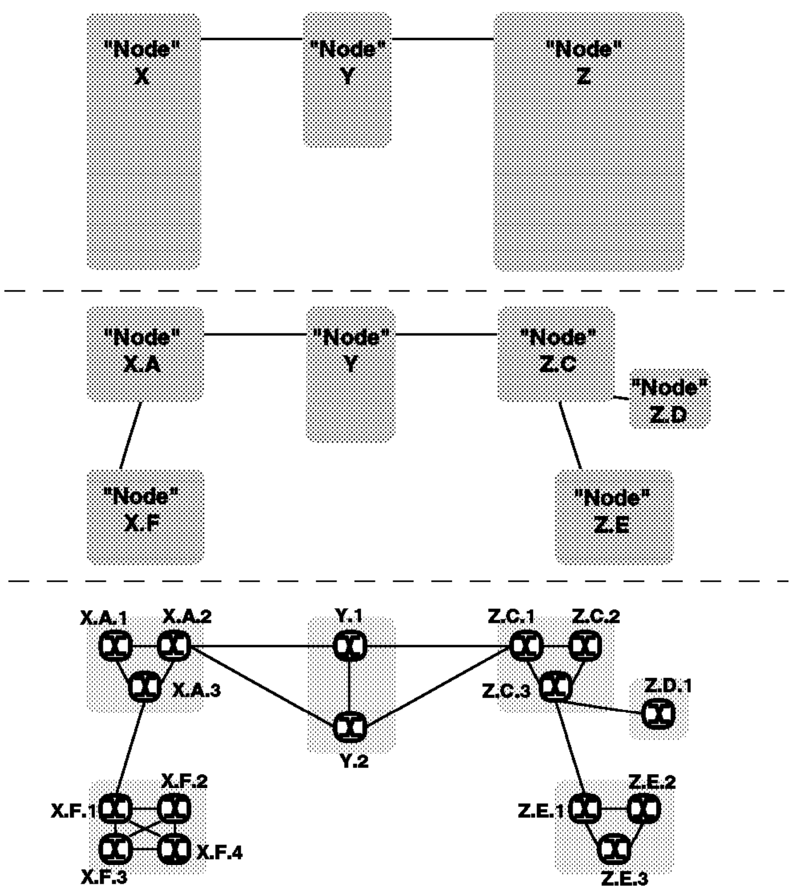

Figure 32. Hierarchical Network Views of PNNI

2.3.6 Data Exchange Interface (DXI)

DXI defines a local interface between a packet-based router and an ATM-capable digital service unit (DSU). The DTE-to-DCE physical interface can be either V.35 or HSSI, so typically the DXI port and the HDLC port of a router can be directly connected. The router then passes packets, encapsulated as defined in RFC1483, to the ATM-capable DSU.

The DSU performs cell level processing, SAR (Segmentation and Reassembly), and virtual connection termination. This allows routers to connect into ATM networks via the UNI interface of the ATM DSU.

2.3.7 The Interim Local Management Interface (ILMI)

Standards for management of ATM networks are still in the formative stage. In the interim period until these are completed, the ATM Forum UNI specification has defined temporary standards based on the use of existing SNMP technology and a new ATM UNI management MIB.

The ILMI protocol uses the above components to allow an ATM end system to get the status and configuration of VP and VC connections available at the UNI. In addition the protocol allows the end system access to other information above the ATM layers, and provides a mechanism for end system registration.

ILMI uses SNMP flows between end system and switch. The flows are encapsulated in AAL5 and a standard channel, VPI=0 VCI=16 is used, although the specification indicates that this should be a programmable option. Message formats are as defined in RFC 1157, that is according to SNMP Version 1, not SNMP Version 2.

ILMI provides access to a range of information that is provided in the MIB defined by the ATM Forum. Examples of the categories of information held within the MIB are as follows:

- Physical Layer
- ATM Layer
- ATM Layer Statistics
- Virtual Path (VP) Connections
- Virtual Channel (VC) Connections
- Address Registration Information

This information is used in a variety of ways including end system address registration, monitoring of interface states, determination of VP and VC availability and many others.

2.3.8 Classical IP

Classical IP is the term used for operation of conventional IP and ARP over an ATM network. The standards for doing this have been defined in RFC 1577.

Classical IP defines the use of IP and ARP in terms that are almost identical to their use on conventional broadcast LANs. ATMARP is used to resolve an IP address to an ATM address, allowing an end system to establish a direct point-to-point connection with the destination end systems. This process is identical to that used by conventional ARP to resolve IP addresses to LAN MAC addresses.

Classical IP operates at the logical IP subnetwork (LIS) level. (Logical because there may be many such subnetworks on an ATM network). Each LIS must have an ATM ARP server to deliver the ARP service to end systems in the network, and all clients must be configured with the ATM address of the ARP server. The ARP server can be a dedicated server or it can be integrated with other functions.

Client end systems in the LIS must be registered with the ARP server at initialization time. The ARP server, therefore, is able to build up a table of ATM addresses and corresponding IP addresses. When queried by a client system it can therefore respond with the destination ATM address for the target IP address. If it does not have the required information for the response, it forwards the ATMARP request to all other clients with which it has a connection. In this way, there is the maximum chance of the IP address being resolved with the LIS.

Once the client end system has the ATM address of the destination end system, it establishes a direct connection over which data frames can be passed. Unlike other architectures, for example LAN emulation, there is no capability for the ARP server to forward data traffic.

The components of a classical IP network and their interaction are shown in Figure 33.

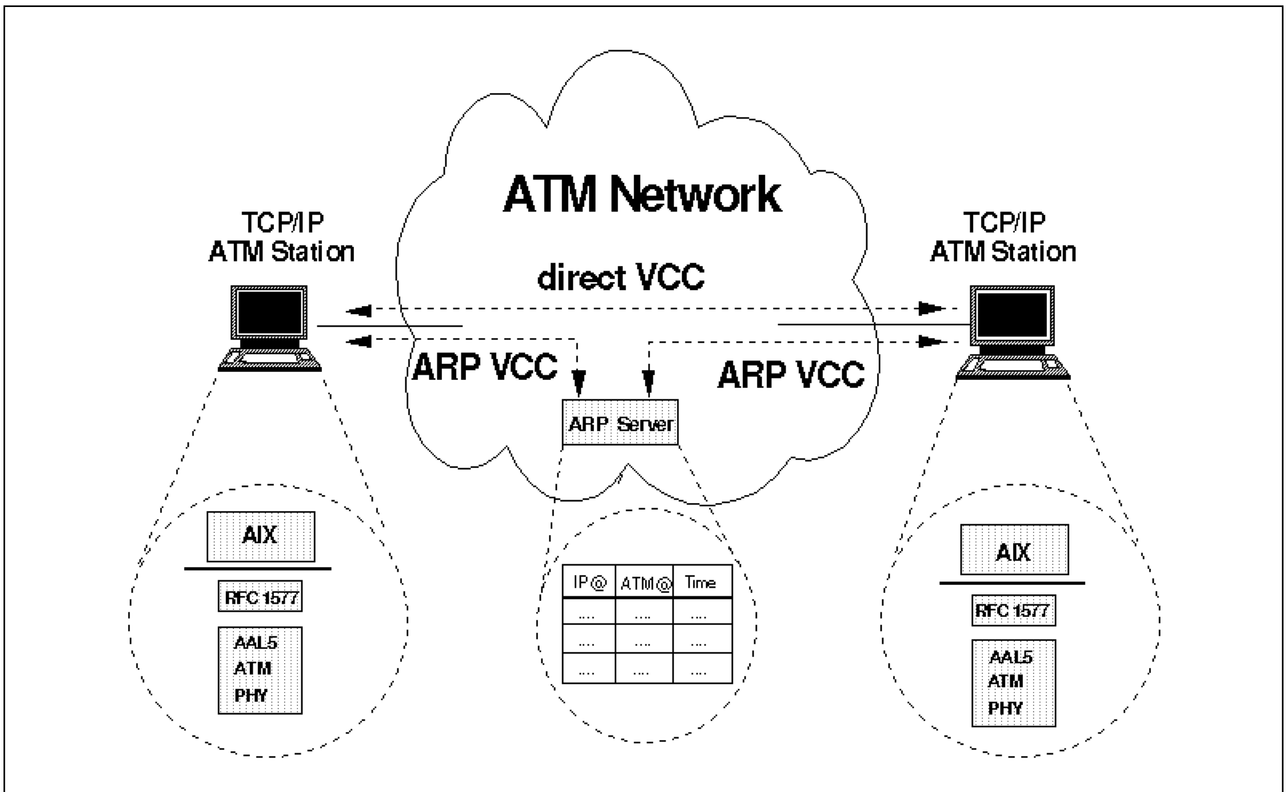

Figure 33. Classical IP Components

2.3.9 LAN Emulation

LAN emulation enables the implementation of emulated LANs over an ATM network. An emulated LAN provides communication between end systems similar to that over a real LAN but using the facilities of the ATM network. LAN emulation software in each end system is provided to allow existing network system software to continue to operate in the new ATM environment. It is the fact that existing applications software can continue to be used that makes LAN emulation so powerful and so important for ATM today prior to the availability of applications using native ATM APIs.

Chapter 3. MPEG-2 over ATM

This chapter describes the concepts of using ATM and MPEG-2 for audiovisual communication services. A brief description of the ITU H.310 videoconferencing standard is also presented.

3.1 ATM and MPEG-2

Asynchronous Transfer Mode (ATM) is an emerging standard for broadband networks that allows a wide range of traffic types, from real-time video to best-effort data, to be multiplexed in a single physical network. A key benefit of ATM technology is its ability to provide quality-of-service (QoS) guarantees to applications. These QoS guarantees are in the form of bounds on end-to-end delay, delay jitter and packet loss rate.

The high-bandwidth support and the multicast capabilities of ATM are also very important characteristics for the transmission of high-quality video.

The multicast support offered by ATM is an extremely useful feature for video distribution solutions. Point-to-multipoint circuits can be easily created and in a standard way, as defined by ATM international standard bodies. This means interoperability and the possibility to implement end-to-end solutions spanning large areas and different service providers.

Endpoint stations can be dynamically added or removed from the multicast tree, through a centralized network management center. A controlled operation of the video distribution is essential to turn it into a viable solution.

Another great ATM feature is that there is no duplication of video traffic on the network links. The ATM network is responsible for the replication of the necessary traffic to the destination points as close as possible to the endpoints. This duplication process is done by the ATM switches in an effective and fast way, through hardware, and with the guarantee of the Quality of Service specified by the user.

This means great savings for the video service provider, since it only pays for the minimum bandwidth needed on the ATM network to deliver the video signal to its customers. As far as the video service provider and its customers are concerned, the ATM network is essentially a *cloud*, transparent to them but responsible for all the operations to guarantee a high-quality service. Figure 34 on page 58 illustrates this mechanism.

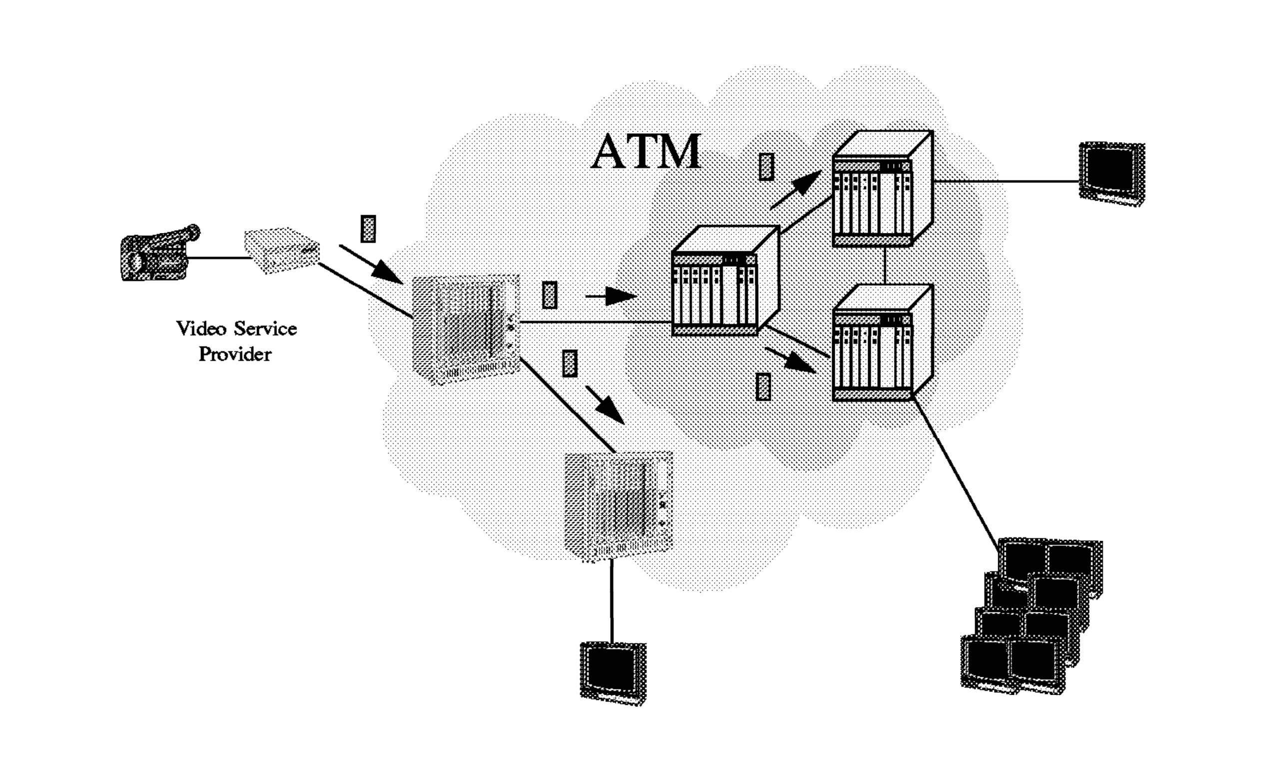

Figure 34. Video Distribution Service Using ATM's Multicast Support

The use of a traffic scheduling algorithm is essential in order to provide QoS guarantees in an ATM network. The function of a scheduling algorithm is to select the packet to be transmitted in the next cycle from all the available packets going through the same output link of a switch. A good scheduling discipline should be able to ensure isolation among the flows so that the desired level of service can be guaranteed to real-time flows even in the presence of other, possibly misbehaving, flows. In addition, the scheduler must also allow best-effort applications to share the available bandwidth.

Several scheduling disciplines have been proposed to achieve this objective. The FIFO (first-in-first-out) scheduling discipline, although the simplest to implement, does not provide any isolation among the various flows and, therefore, cannot offer deterministic bandwidth guarantees. Many algorithms that provide bandwidth guarantees to individual sessions are known. Many of these algorithms are also capable of providing deterministic delay guarantees when the burstiness of the session traffic is bounded, as in the case of the output of a leaky-bucket shaper.

MPEG-2 is the emerging standard for audio and video compression. Being capable of exploiting both spatial and temporal redundancies, it achieves high compression ratios and can encode a video or audio source to almost any level of quality.

The MPEG-2 standard offers two ways to multiplex elementary audio, video or private streams to form a program: the MPEG-2 Program Stream and the MPEG-2 Transport Stream formats. The MPEG-2 Transport Stream is the approach suggested for transporting MPEG-2 over noisy environments, such as an ATM network. Using explicit timestamps (called Program Clock References or PCRs in MPEG-2 terminology), MPEG-2 Transport Streams ensure synchronization and continuity, and provide ways to facilitate the clock recovery at the decoder end.

3.1.1 ATM Adaptation Layers

The MPEG-2 standard does not specify how an MPEG-2 Transport Stream is transported over a communication network. In the ATM world, the adaptation layer is responsible for making the network behavior transparent to the application. It is divided into two sublayers: the Segmentation and Reassembly (SAR) and the Convergence Sublayers (CS). The SAR is responsible for the segmentation of the outgoing Protocol Data Units (PDU) into ATM cells and the reassembly of ATM cells back into PDUs.

There are four types of adaptation layers currently defined for ATM networks: AAL1, AAL2, AAL3/4 and AAL5. Each is designed for supporting specific services and have different functionalities. The selection of a suitable adaptation layer for transporting MPEG-2 over ATM needs to take into account the specific requirements of MPEG-2 Transport Streams, such as jitter removal, error detection and/or correction, end-to-end delay minimization for real-time applications and support of both CBR and VBR applications.

3.1.1.1 Transport over AAL1

AAL1 was designed to support circuit emulation over ATM networks. It is ideally suited for transporting Consultant Bit Rate (CBR) traffic since it provides constant delay through the network using dejittering mechanisms at the destination. AAL1 provides two ways to synchronize the clocks and deliver a jitter-free clock at the receiver depending on the status of the CBR service clock:

AAL1 also offers the option of Forward Error Correction (FEC) which can hide the effects of cell losses in the network from the application.

AAL1 is currently the natural choice for transporting CBR traffic over ATM since the traffic has a constant bit rate and needs constant end-to-end delay. However, there are a number of disadvantages using AAL1 for MPEG-2 transport:

- AAL1 cannot be used to carry Variable Bit Rate (VBR) MPEG-2 Transport Streams which are likely to be dominant in the future.
- The SRTS technique cannot be used if a common network clock is not available for use as the reference clock. That is, AAL1 cannot be used in nationwide networks consisting of several carriers and unsynchronized clocks.
- The adaptive method to recover the clock requires a PLL to determine how the buffer at the decoder end is being emptied. Since a PLL is needed for recovery of the MPEG-2 system clock, the former becomes redundant.
- Since signaling is being done under AAL5, ATM network interfaces will need to support both types of adaptation layers (AAL1 and AAL5).

3.1.1.2 Transport over AAL5

AAL5 was designed for transporting data traffic with no real-time constraints over ATM.

Errors in the transmission of the data units can be detected, but cannot be corrected through a forward error correction method as the one used by AAL1.

AAL5 has several advantages over other alternatives:

- AAL5 is currently the most commonly used adaptation layer in the industry. AAL5 is being used for encapsulating UNI 4.0 signaling messages and, in most cases, to carry best effort traffic through the ATM network.
- Using a null CS, hardware support from the network can be minimized and thus the complexity is moved to the service layer.
- AAL5 can support VBR MPEG-2 traffic in the future.

However, there are disadvantages also, such as the lack of forward error correction. AAL5 only supports error detection and, in the case of an error, may discard the entire service data unit, which usually amplifies the effect of the error in audiovisual applications. However, since congestion losses in an ATM network are likely to occur in bursts, the effectiveness of forward error correction may be limited in any case.

3.2 The H.310 Standard

The combination of MPEG-2 and ATM makes possible broadband audiovisual communication services. The ITU has determined an international standard for these broadband audiovisual services. The standard, H.310, specifically identifies MPEG-2 and ATM as the basic elements of these services.

The transmission format of H.310, called H.222.1, is defined as the MPEG-2 systems layer transport stream standard. Material in this format can be stored and played back readily. Only H.310 provides this capability.

Audiovisual services imply much more than simple transmission of video across a network. To provide a reliable end-to-end service, networked video equipment must be able to exchange information with each other using a multimedia control protocol. The H.310 standard references H.245, a multimedia control protocol that allows peer H.310-aware audiovisual services equipment to set up and maintain a broadband audiovisual session.

The MPEG-2 transport stream is designed to carry multiple audio and video streams across reliable data networks. The standards-setting body of ATM, the ATM Forum, has recommended that ATM Adaptation Layer 5 (AAL5) be used to transmit MPEG-2 transport stream data. Figure 35 on page 61 illustrates this concept.

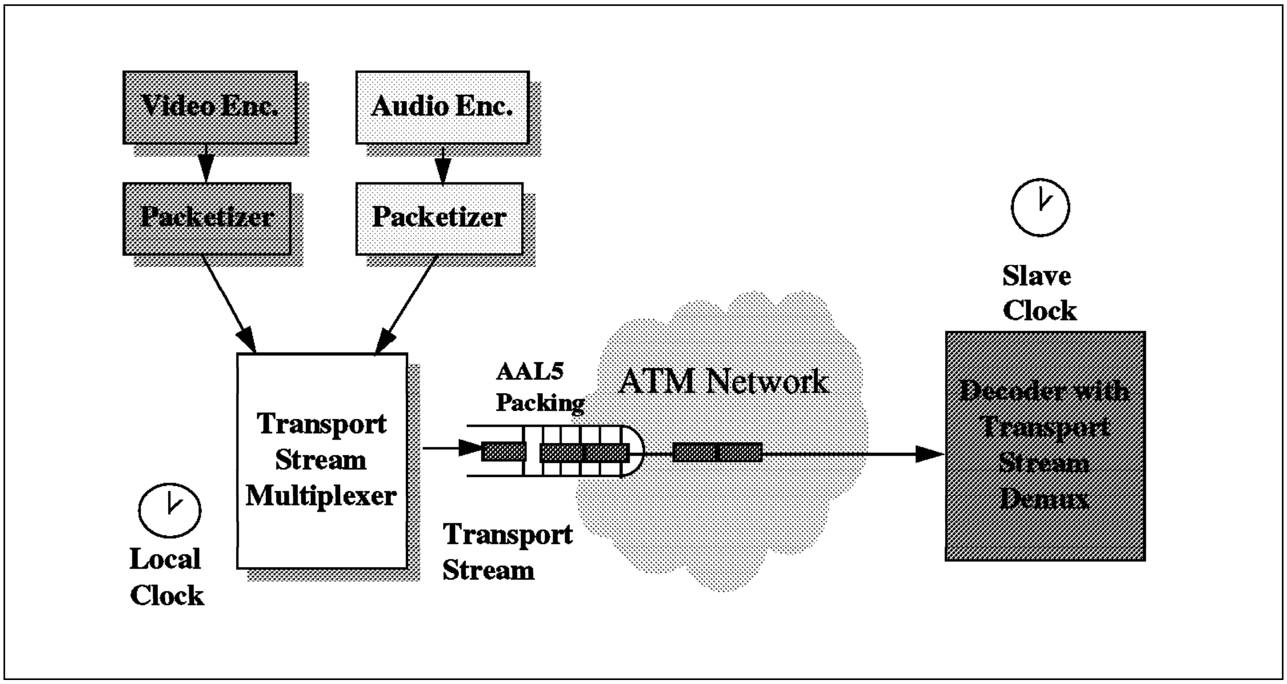

Figure 35. MPEG-2 over ATM

As networked video becomes more mainstream, the total network resources dedicated to high-bandwidth networked video solutions are likely to become a significant portion of network traffic. ATM networks provide on-demand bandwidth at the requested Quality of Service. Thus, ATM-based networked video systems enable dynamic networked video connections. However, even on ATM networks, the significant bandwidth consumed by MPEG-2 video must be managed carefully.

To simplify the operations of a video transmission system, the service can be managed through a network-centric control and management model. This model provides a single point of management for all network resources through schedulers and system management facilities.

Unlike M-JPEG, H.310 offers interoperability between vendors, transmission over ATM, and very high-quality video using less bandwidth. MPEG-2 also allows recording of audiovisual material, a feature not available with H.320 or M-JPEG.

MPEG-2 over ATM videoconferencing offers a unique solution as today's videoconferencing choices are either desktop quality (H.320) or expensive and not interoperable (M-JPEG over DS-3).

Chapter 4. IBM Solutions for Networked Video

The recent progress in audio and video signal processing has made possible the development of numerous applications involving multimedia information. In this chapter we present the integrated, managed and standards-based solutions offered by IBM for networked video applications.

4.1 Introduction

Video transmission systems used in the television industry connect network video-production facilities with affiliates, advertisers, CATV headends, field production units and other studios across the world.

Today′s transmission systems typically carry composite video in an analog or digitized form. Short-distance transmissions use terrestrial microwave or dedicated-fiber links. Cross-country or transoceanic hookups require satellite transponder time at rates exceeding $1000/hour. The increasing demand for video transmission, coupled with capacity constraints and the limited expansion rate of satellite services, will push these rates even higher in the near future.

The high-quality video compression and transmission provided by MPEG-2 over Asynchronous Transfer Mode (ATM) networks naturally lends itself to video transmission applications for the broadcast industry.

IBM′s networked video technologies combine IBM′s MPEG-2 encoding chips with its leading ATM transmission and switching experience. These technologies enable the transition from today′s analog transmission systems to fully digital networked video infrastructures. IBM′s networked video technologies open a new era for video transmission in terms of cost, capacity, manageability, and flexibility.

Terrestrial ATM networks now span much of the United States. With networked video technologies, broadcast-quality video can be transmitted at rates less than 15 Mbps over ATM. A typical ATM fiber link can carry from 155 Mbps to 4.8 Gbps. This translates to 10-300 connections on a single fiber. With wavelength-division multiplexing, up to 160-4800 connections can be transmitted on the same fiber. With such transmission densities, we expect most short-distance transmissions to move to ATM networks and use MPEG-2 compression for significantly lower cost and higher capacities, in the near future.

Long-distance transmissions also benefit from ATM. Most long-distance transmissions have multiple destinations. For example, when a basketball game is played, there will be at least three recipients of the video signal from the game site: the home city, a network headquarters city such as New York, and at least one satellite uplink site. When video signals are broadcast to relatively few locations, ATM′s multicast facilities can optimize the link use over the network. Depending on the topology of the network, it is conceivable that the net bandwidth use (sum of total bandwidth consumed on links in the network) will not greatly exceed the use of a single transmission and will be far less than that of multiple point-to-point links used today.

In today′s broadcast systems, the complex mix of analog production equipment, transmission equipment, and satellite connections is not covered under a uniform, system-management umbrella because of the heterogeneity of the equipment and services. Traditional data and voice networks have provided comprehensive network-management facilities for many years. The merger of well-understood network management with networked video technologies will offer the best of both worlds: seamlessly integrated and managed networked-video systems for video transmission.

Even a managed network requires interoperability between endpoints. The H.310 Broadband Audiovisual Services standard from the International Telecommunication Union (ITU) defines profiles that allow MPEG-2 over ATM equipment and services to interoperate. IBM′s networked video technologies implement H.310-compatible broadband, audiovisual services.

4.1.1 Interactive-Video Solutions

The field of education uses a form of videoconferencing for distance learning. Distance learning uses conferencing to allow lecturers to reach students at remote locations and permit full interactivity with the instructor. The digitized program may also be stored for later retrieval at a student′s convenience.

The business world has a need for time-critical, information-exchange meetings. In a fast-paced business environment, there is an increasing need for face-to-face meetings with counterparts in other parts of the world. With constraints on travel, money and time, broadcast-quality, full-motion, interactive videoconferencing provides a cost-effective way to exchange information without leaving the confines of the conference room.

The medical community can use the videoconferencing solution to exchange data and voice for collaborative efforts in remote locations. Exchange of information and expertise are possible through real-time broadcast transmission.

Television can use the new networking offerings from IBM for real-time broadcast transmission, such as that required by sporting events, trade shows and entertainment. Network-to-station just-in-time advertising is also possible.

Business and education have requirements for video on demand. The capability to store audiovisual material and to provide scheduling guarantees that the material and the bandwidth will be there when they′re needed.

An integrated videoconferencing solution such as the one illustrated on Figure 36 on page 65, can be implemented with the wide range of IBM offerings in the video and ATM areas. ATM switches, such as the 8260 Nways Multiprotocol Switching Hub, 8285 Nways ATM Workgroup Switch Expansion Unit and the 2220 Nways BroadBand Switch are in operation throughout the world.

Audiovisual solutions include the Video Access Node, the Video Distribution Module and the MediaSTREAMER. The following sections describe these solutions in more detail.

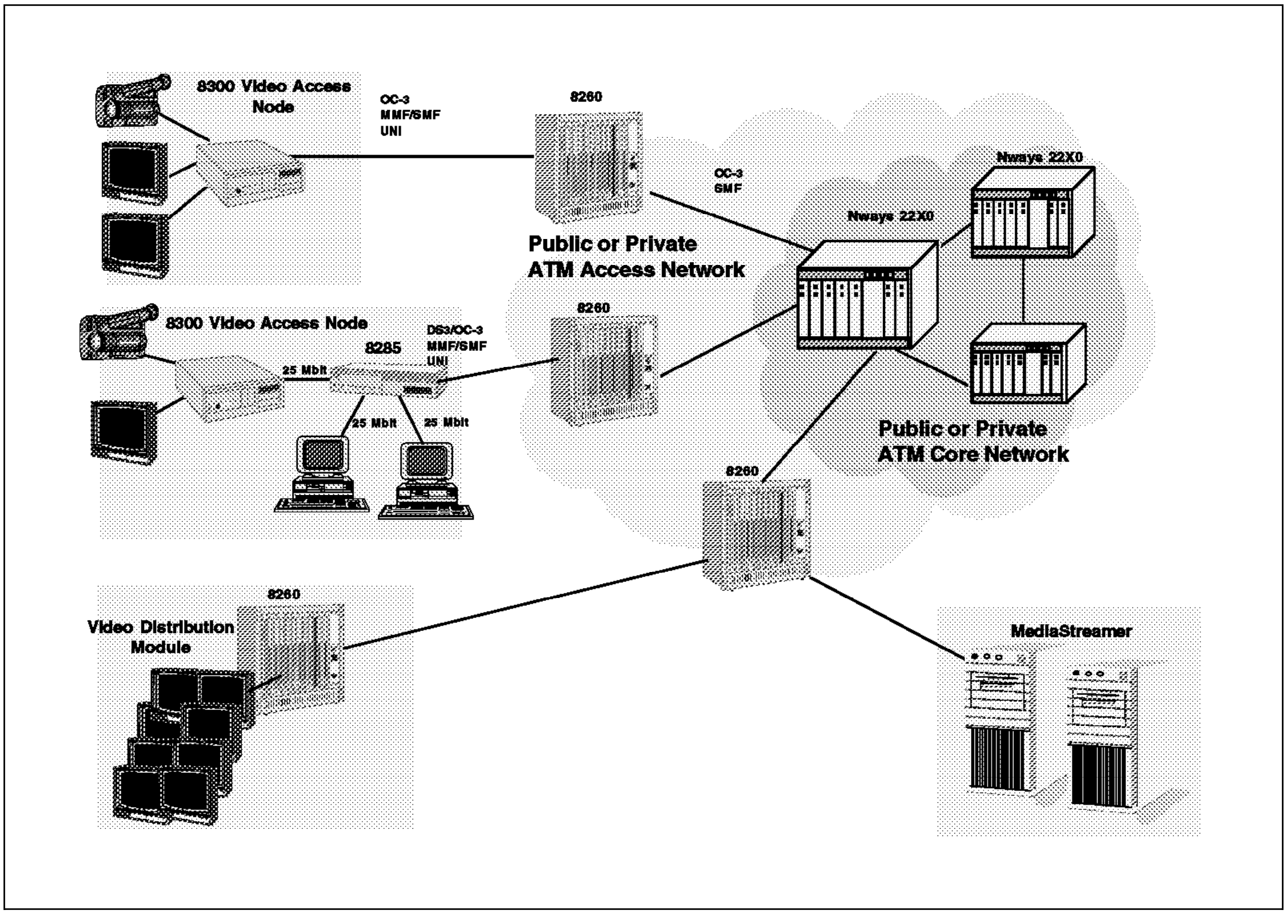

Figure 36. Videoconferencing Solution

4.2 Video Access Node (VaN)

The Video Access Node (VaN) provides real-time, interactive delivery of broadcast-quality video over ATM networks. A typical VaN application converts analog video and audio signals into a digital form, transmits the signal over ATM to a remote endpoint, decodes the signal and displays the signal, for example, on a television monitor.

The VaN provides:

- Cost-effective, broadcast-quality video compression and transmission
- Simultaneous encoding and decoding of video and audio data
- Standards-based, synchronized audio and video
- Merger of video with other traffic on the ATM network
- Recording of audiovisual material for storage and playback
- The capability to customize networked video applications

The VaN adheres to the following standards to achieve high-quality video compression and transmission over ATM networks:

- NTSC and PAL formats.
- The MPEG-2 standard for digitizing, compressing, and multiplexing video.
- The MPEG-1 standard for audio encoding.

- A subset of the ITU H.310 Broadband Audiovisual Communications Systems and Terminals standard for transmission of MPEG-2 data over ATM networks.
- ATM Forum Video-on-Demand (VoD) terminal standard.
- The VaN conforms to UNI 3.1 and takes full advantage of ATM′s capabilities.

The Video Access Node is a programmable platform. For application development, IBM provides an Audiovisual Services application programming interface (AVS API). This enables customers and value added remarketers (VARs) to develop different and customized networked video applications.

Applications for the Video Access Node include two-way video communication such as distance learning, business video conferencing, and medical collaboration (tele-medicine). The VaN also enables video transmission applications such as terrestrial television signal transmission for station to venue signal transmissions.

4.2.1 Broadband Audiovisual Communication

A *broadband audiovisual communication system* enables the transmission of audiovisual data from one audiovisual terminal to one or more other audiovisual terminals over an ATM network.

An *audiovisual terminal* is a terminal that can accept audio and video data from an input device and present audio and video data to an output device over an ATM network. A *broadband audiovisual terminal*, additionally, can send, receive, or both send and receive audiovisual data through an ATM network.

Figure 37 on page 67 shows the components of a broadband audiovisual communication system that includes the VaN.

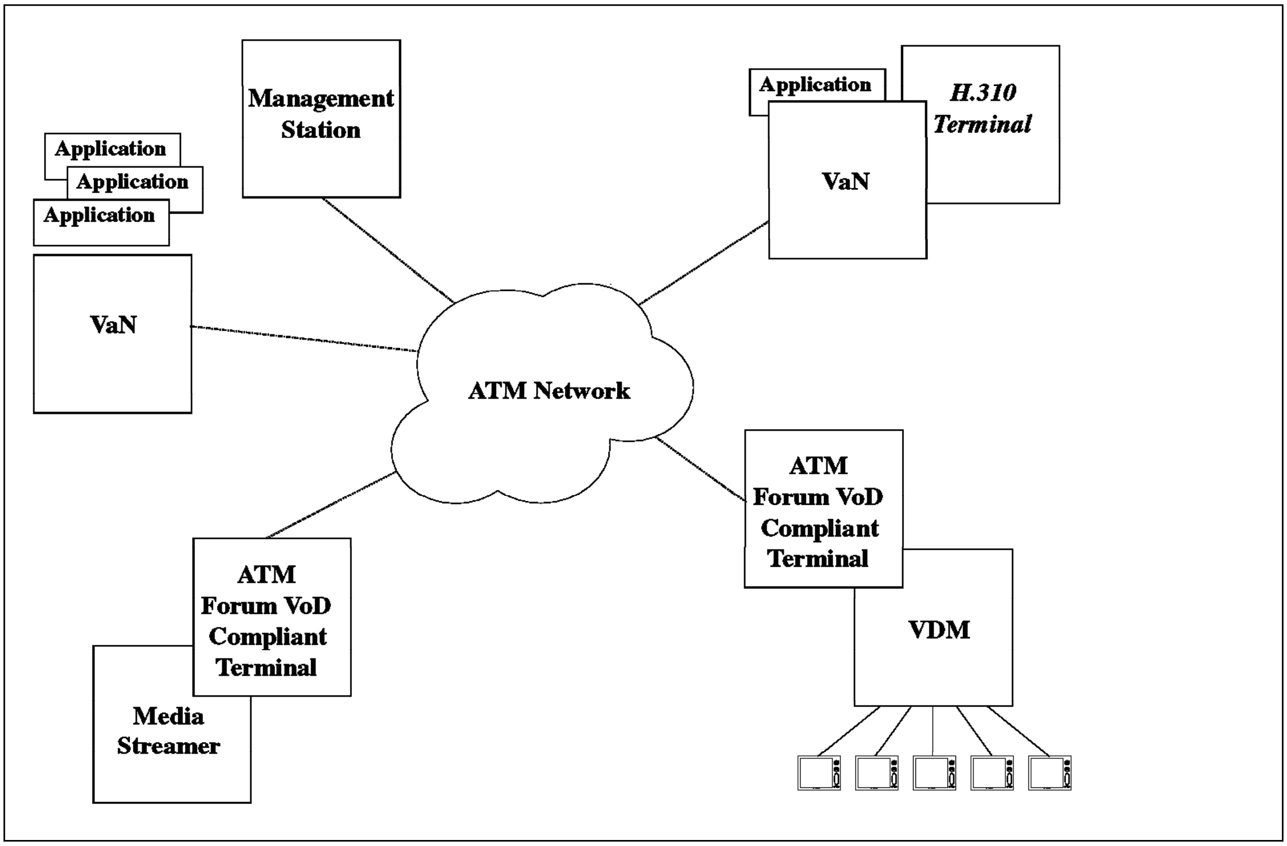

Figure 37. VaN Network

The VaN is based on a subset of the H.310 specification, which describes the hardware and software requirements for exchanging control and audiovisual data between compliant terminals. Because the terminal functions implemented in the VaN are compliant with H.310, the VaN is referred to as an H.310 terminal.

The VaN can send audiovisual data to and receive audiovisual data from other VaNs; the VaN also can send audiovisual data to and receive audiovisual data from terminals that are compliant with the ATM Forum Video on Demand (VoD) standard. Terminals that are compliant with ATM Forum VoD do not require the exchange of control information prior to transmitting data, but do support both the ITU J.82 standard and the ATM Forum standard for VoD. Examples of such terminals are IBM's VDM and MediaSTREAMER products.

The VaN uses SNMP over TCP/IP to communicate management information to a network management station. TCP/IP communications can be connected through the ATM network or through a local area network.

VaN software includes an application programming interface (API) for developing applications.

4.2.2 VaN Hardware Description

The VaN hardware is based on a PCI and EISA bus backplane with multiple adapters. Depending on your requirements, you can populate the VaN with a variety of configurations. The VaN has one EISA, one PCI, and four shared PCI/EISA slots.

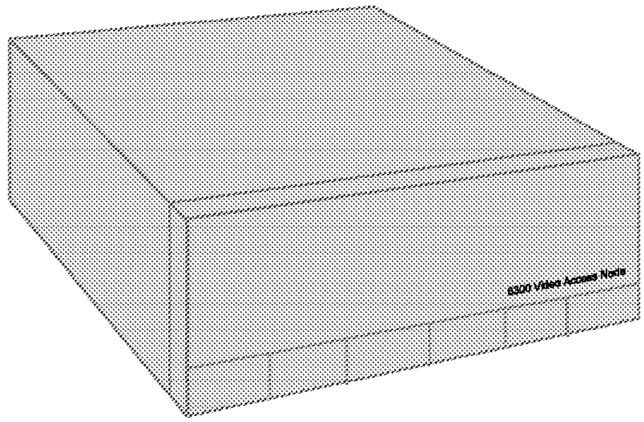

Figure 38. The Video Access Node

Base hardware is considered to be that which ships from the factory on or with the VaN unless other hardware is specified in the order. Features are usually optional, and each feature is identified by a four-digit feature code.

The base hardware for the VaN includes:

- The system board and processor module (a Pentium 200-MHz processor)
- A 2-GB hard file
- A CD-ROM drive
- 32 MB of memory
- A 3.5-inch diskette drive
- A keyboard and mouse
- An external modem in the U.S., Canada, and some other countries

The system board contains zero-insertion-force (ZIF) sockets for the processors. The one EISA, one PCI, and four shared PCI/EISA slots on the system board provide connections for:

- MPEG-2 Encoder Adapters
- MPEG-2 Decoder Adapters
- An ATM adapter
- A Token-Ring Adapter

The system board also contains an integrated Ethernet network connection with RJ-45 and DB-15 AUI connectors.

The hard file is a direct access storage device that is preloaded with OS/2 and the AVS software.

The 8x-speed CD-ROM drive has a small computer system interface (SCSI).

A 32-MB Dual Inline Memory Module, or DIMM, is plugged into a connector on the processor board.

A modem is required for the remote service of this product. A 28.8 kbps, V.34, asynchronous, Hayes-compatible external modem is provided with the VaN in the U.S., Canada, and some other countries.

The VaN supports many hardware features. These features, and combinations of them, enable the VaN to support varied application scenarios. Features can be factory- or field-installed.

A brief description of the supported features is presented below:

The MPEG-2 Encoder Adapter is a two-slot PCI adapter that provides video and audio encoding. A maximum of two Encoder Adapters are supported by the VaN. You must order one of the two type of audio cables with each Encoder Adapter.

The Balanced Audio Cable is a 3-meter (9-ft, 10-inch) balanced audio cable for the MPEG-2 Encoder Adapter. It has female XLR connectors.

The Unbalanced Audio Cable is a 3-meter (9-ft, 10-inch) unbalanced audio cable for the MPEG-2 Encoder Adapter. It has two male RCA connectors.

The MPEG-2 Decoder Adapter is a single-slot PCI adapter that provides MPEG-2 video decoding. A maximum of four Decoder Adapters are supported.

The 155-Mbps Multimode Fiber ATM Adapter provides ATM connectivity over multimode fiber. One PCI ATM adapter is supported and required in the VaN.

The 155-Mbps UTP5/STP ATM Adapter provides ATM connectivity over UTP5 or STP copper wire. One PCI ATM adapter is supported and required in the VaN.

The 25-Mbps UTP/STP ATM Adapter provides ATM connectivity over UTP category 3 to 5 or STP copper wire. One PCI ATM adapter is supported and required in the VaN.

The Token-Ring Adapter provides connectivity to 16/4-Mbps and full-duplex token-ring networks. This optional Token-Ring Adapter and the Ethernet connection that is integrated into the base machine allow you to use a traditional LAN environment for applications communication. A maximum of one Token-Ring Adapter is supported. Both EISA and PCI adapters are supported.

The 15-inch Multisync Monitor is a color monitor providing a 13.5-inch (342 mm) diagonal viewable image. A monitor must be available to the VaN for remote service. Either the 15-inch Multisync Monitor or a compatible, customer-supplied equivalent must be provided.

The Rack-Mount Kit allows you to mount the VaN in an EIA 19-inch rack.

A 4-GB Hard File can be ordered to supplement the 2-GB hard file included with the VaN. You would consider this option if you want to store limited amounts of video on the VaN.

The following list contains guidelines for configuring the VaN hardware:

- One ATM adapter (PCI form factor) is supported and required.
- Use the four remaining PCI slots for Encoder Adapters and Decoder Adapters. A maximum of four slots can be used for Encoder Adapters and Decoder Adapters. Because the Encoder Adapter is a two-slot adapter, a fully populated VaN will accommodate these combinations:
  - One Encoder Adapter and two Decoder Adapters
  - Two Encoder Adapters

- Four Decoder Adapters

You also can combine Encoder Adapters and Decoder Adapters in combinations that do not fully populate the VaN, such as:

- One Encoder Adapter
- One Encoder Adapter and one Decoder Adapter
- One to three Decoder Adapters

- One optional Token-Ring Adapter (PCI or ISA form factor) is supported. The selection of the form factor for this adapter is based on available slots after designation of Encoder Adapters and Decoder Adapters. ISA is the default form factor.

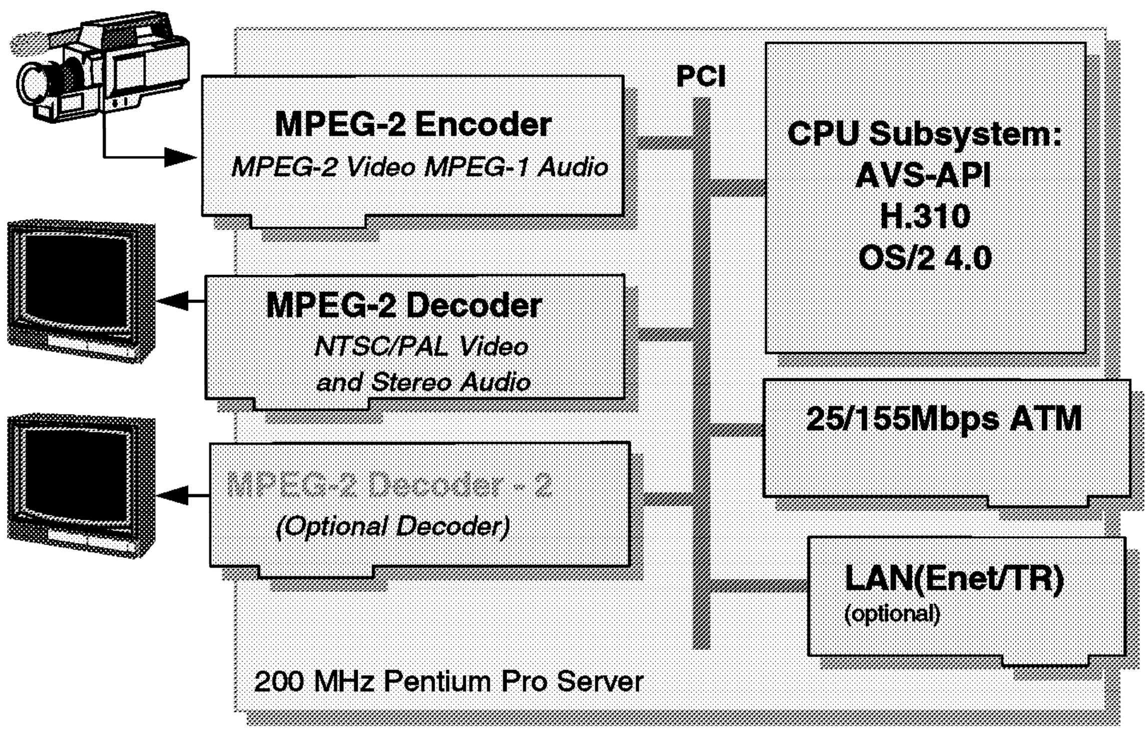

Figure 39. VaN Hardware

4.2.3 VaN Software Description

This section describes the components of the VaN software and the H.310 specification upon which the software is based.

The VaN hard file is preloaded at the factory with the VaN software, which includes:

- IBM Operating System/2 (OS/2) V4.0 and IBM TCP/IP V4.0 with ATM extensions.
- AVS software allowing you to configure and operate the VaN. The AVS software also provides an API for developing your own applications.

The AVS software allows you to configure hardware, use a command line interface for such tasks as starting audiovisual transmission, and provides a set of functions that allows applications to:

- Set up unidirectional, bidirectional, and multicast audiovisual transmissions over ATM
- Control encoding and transmission parameters
- Query statistics and status of VaN resources

Applications will run on OS/2 V4.0. Software developers must use a customer-supplied C or C++ compiler that supports OS/2 V4.0. In a typical networked system, a minor software development effort will provide the client stub on the VaN that will allow the use of existing network service managers with the VaN.

4.2.3.1 H.310 Compliance

Figure 40 illustrates the VaN's implementation of the H.310 protocol model, which is a subset of the overall standard.

Refer to the ITU H.310 Broadband Audiovisual Communications Systems and Terminals standard for complete information about the H.310 standard.

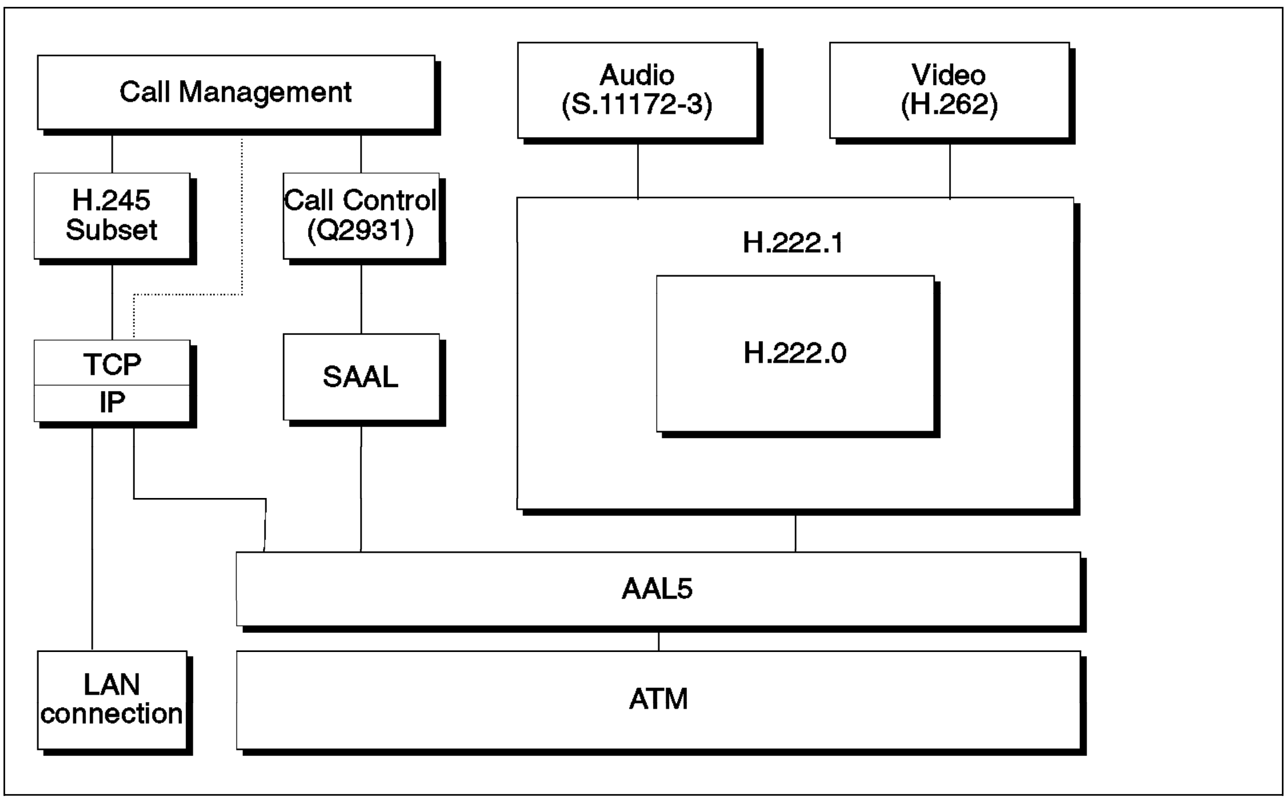

Figure 40. VaN H.310 Protocol Submodel

In the VaN's H.310 implementation:

- Control messages (via the H.245 protocol) are exchanged over TCP/IP instead of OSI-reliable transport as specified by H.310.
- The VaN does not support interoperability with any other audiovisual terminal standards, such as H.320 or T.120.
- H.310 Video Frame Synchronous commands and indications are not supported.
- A small number of H.245 elements that are required by H.310 are not supported. The H.245 protocols supported by the VaN are:
  - Master/slave determination
  - Capabilities exchange
  - Logical channel signaling (unidirectional and bidirectional)
  - Close logical channel signaling
  - Round-trip delay determination
  - Maintenance loop signaling
  - A small number of commands and indications (SendTerminalCapabilitySet, EndSessionCommand, FunctionNotSupported, and NewATMVCIndication)

  All other H.245 protocols are unsupported.

These H.310 subcomponents are described in the following sections in terms of their H.310 composition and their implementation of the subcomponent in the VaN:

- H.222.1 and H.222.0
- Call control and SAAL
- ATM and AAL5

H.222.0 and H.222.1: H.222.0 (ISO/IEC 13818-1) describes MPEG-2, which is a method for compressing and multiplexing audiovisual data. H.222.1 describes how to adapt MPEG-2 streams for transmission across an ATM network. H.310 specifies H.222.0 and H.222.1 as its multiplexing systems and network adaptation layers, respectively.

The VaN supports the following portions of these standards:

- MPEG-2 transport streams but not program streams
- ATM adaptation level 5 (AAL5) (not AAL1)
- A minimum of 2 and a maximum of 348 transport stream packets per AAL5 protocol data unit (PDU)

Call Control and SAAL: The ATM components include Q.2931 signaling and Q.2100 Signaling ATM Adaptation Layer (SAAL) reliable communications to carry the signaling messages between the H.310 terminal and the network switch in the ATM network. The VaN does not use these standards for control connections; all of the call control messaging in the VaN is carried out over a TCP/IP connection. Q.2931 is used by the VaN to establish audiovisual sessions.

ATM and AAL5: The H.310 standard specifies an ATM physical layer as well as AAL1 and AAL5 as the transmission medium for the multiplexed data. AAL5 is used also to transmit and receive commands issued by the H.245 protocol stack. The VaN supports only AAL5.

4.2.3.2 Support in Addition to the H.310 Standard

The VaN goes beyond H.310 by providing support for transmission and reception of raw MPEG-2 transport streams between the VaN and ATM Forum VoD terminals.

H.310 implies that a terminal will have at most one input and one output device. The concept of the logical terminal is used by the VaN to allow multiple instances of similar devices (more than one encoder and more than one decoder), while still fitting the H.310 model.

Figure 41 shows a simplified view of the traffic flows and the software involved in a VaN-to-VaN connection.

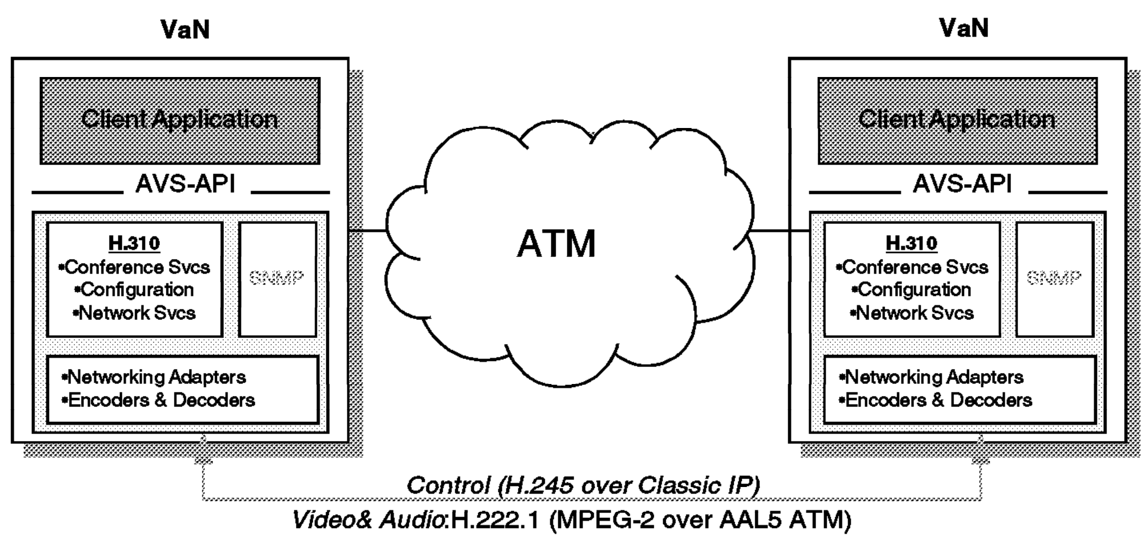

Figure 41. VaN Components Involved in a Videoconferencing Scenario

4.2.4 Network Management

The VaN is a remotely manageable product. It supports:

- SNMP management

  The VaN provides an SNMP management information base (MIB) with settable alerts, traps, and thresholds.

- System monitoring and operation

  A command interface is available for activation, deactivation, and testing of the VaN.

- Remote software installation and upgrade

  The VaN is enabled for configuration, installation, and distribution (CID-enabled), permitting the software to be remotely installed and upgraded.

You may use a network-centric service manager to remotely control the establishment of audiovisual session on your VaNs. You may run a local application through a Web browser, for instance. This feature provides a very powerful and flexible interface for end users or the network managers. Figure 42 on page 74 illustrates this concept.

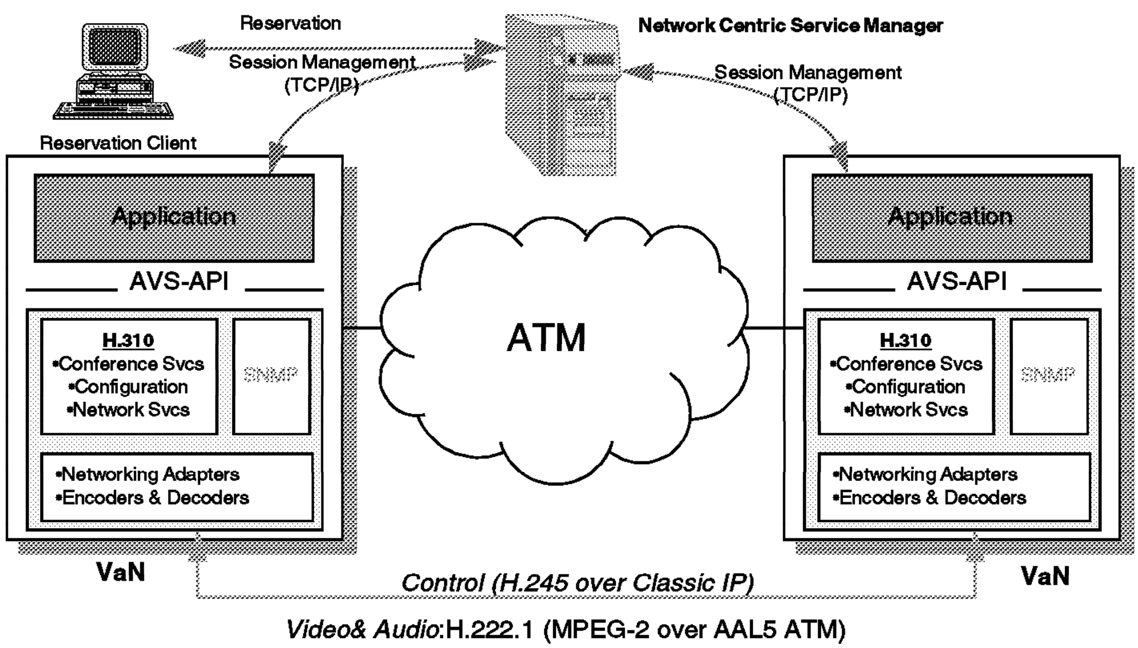

Figure 42. Remote VaN Control with Network Centric Service Manager

4.3 Video Distribution Module (VDM)

The Video Distribution Module (VDM) is a two-slot module for the 8260 Nways Multiprotocol Switching Hub and the 8285 Nways ATM Workgroup Switch Expansion Unit. Figure 43 on page 75 shows the main hardware components of the VDM.

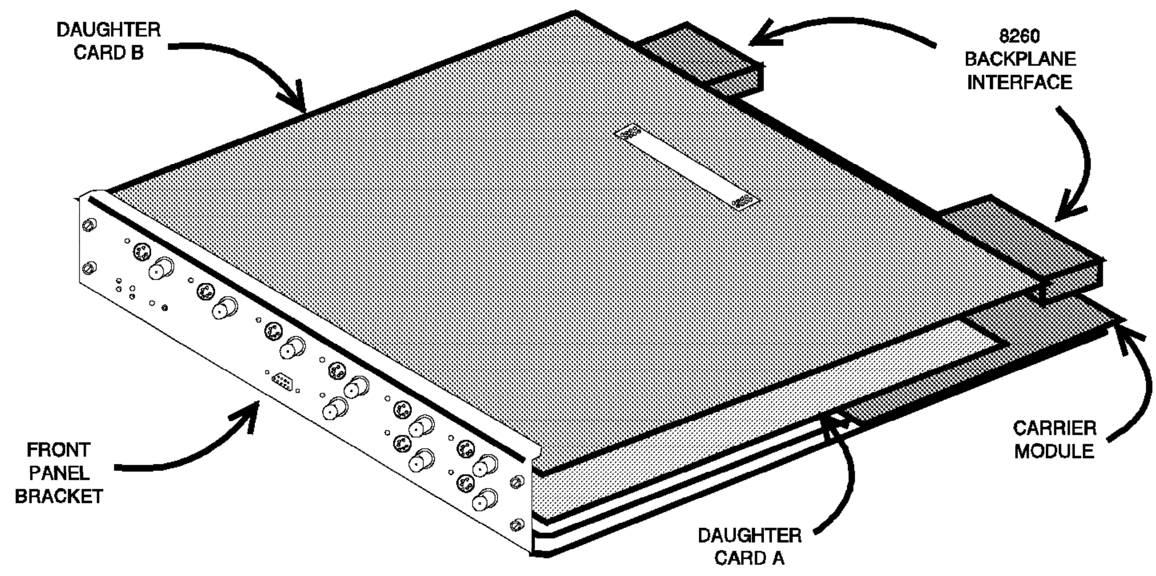

Figure 43. The Video Distribution Module

A video source, such as a video server or the IBM 8300 Video Access Node, can establish a connection with one or more of the VDM ports and transmit an MPEG-2 stream across the ATM network. The VDM receives the MPEG-2 transport streams from the ATM switch in the 8260/8285, decompresses and decodes the video information, and converts it into baseband analog outputs. The VDM contains eight independently addressable MPEG-2 decoder ports. Each port provides separate analog video in NTSC or PAL format and CD-quality stereo audio outputs.

The VDM also has the ability to receive closed caption information in the MPEG-2 transport stream and reinsert that information into the Vertical Blanking Interval (VBI) of the outgoing analog video signal. Support for VBI data is a key requirement of the broadcast industry, and the ability to handle closed caption data is the first stage of VBI support.

Figure 44 on page 76 shows a typical video application using the VDM.

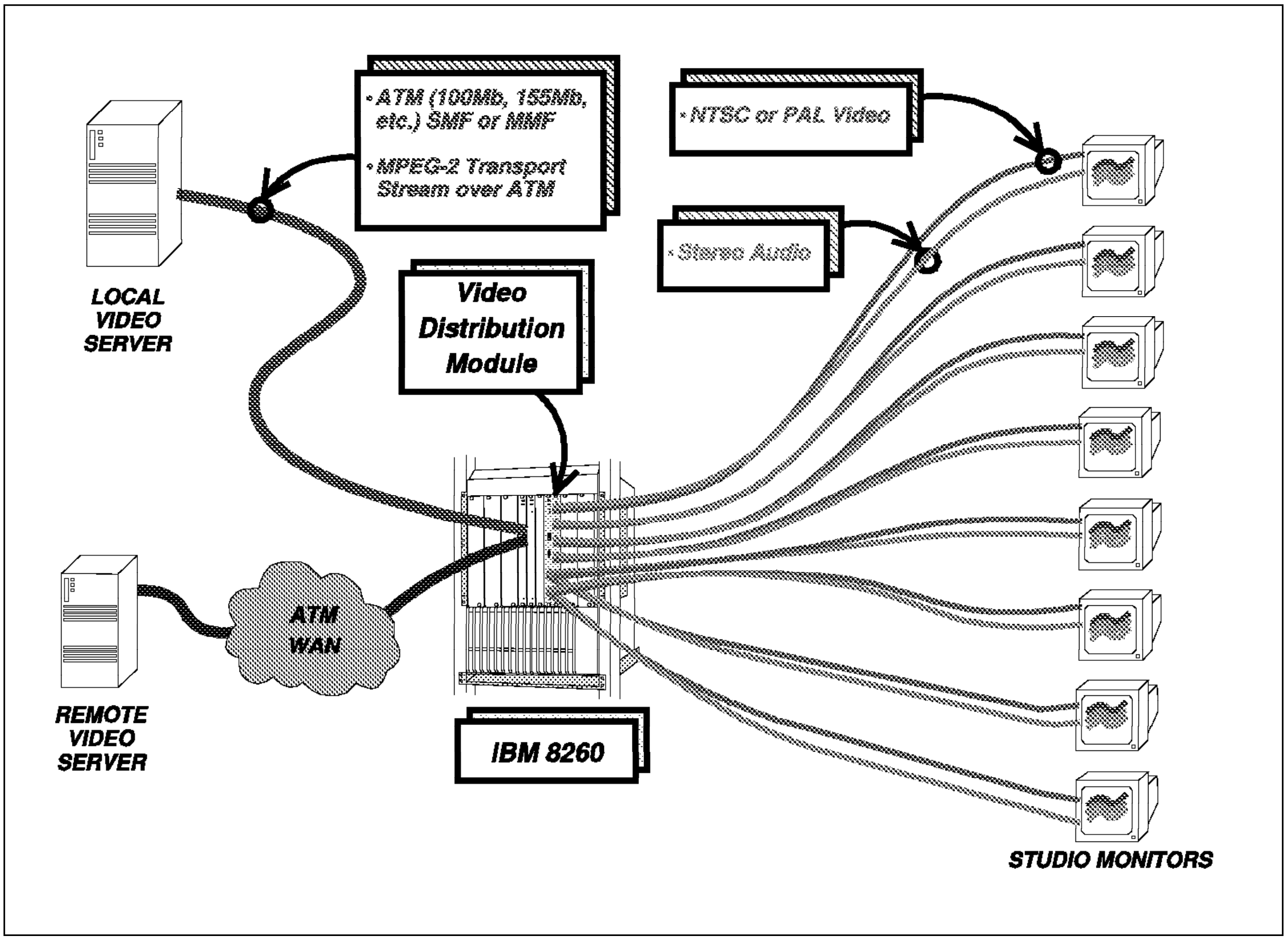

Figure 44. Video over ATM

4.3.1 Main Characteristics

The VDM connects to the ATM backplane through the ATM Carrier Module. The VDM appears as a switched resource of the 8260/8285 and functions as an ATM end node, supporting ATM UNI 3.1 signaling.

The VDM module has the following characteristics:

- Aggregates throughput of 125 Mbps over eight ports (15.4 Mbps maximum per port).
- Separates analog composite video (NTSC or PAL) and stereo audio outputs on each port.
- Power requirement: 62.5 W at +5 V dc.
- Physical interfaces:
  - Eight video ports - BNC
  - Eight audio ports - 5 pin DIN
  - One EIA-232 console port - 9 pin D-shell
  - One genlock port - BNC
  - One ATM backplane connector
  - One trichannel backplane connector
- Cables provided:
  - Eight video cables - 3 m (9 feet, 10 inches)

RG59/U (75 ohms impedance
  - Eight audio cables - 3 m (9 feet, 10 inches) unbalanced

- Cables not provided:
  - EIA 232 console port cable set - IBM P/Ns 59G0278/58F2861/58G4422 (These parts are provided with the 8260 Hub)
  - Genlock - RG59/U (75 ohms impedance)

- Cable lengths supported:

  Video 75 m (246 feet)
  Audio 75 m (246 feet)
  EIA 232 15 m (49 feet)
  Genlock 75 m (246 feet)

- Additional features:
  - Maximum number of VDMs per hub:

    8260 Model A17 – 7
    8260 Model A10 – 3
    8285 Model 00E – 1

  - The VDM module can be hot-plugged while the hub is in operation without disturbing data traffic on other modules in the hub.
  - The VDM supports permanent virtual circuit (PVC) or switched virtual circuit (SVC) mode ATM connections.
  - The VDM microcode can be upgraded and configured via the console port.

4.3.2 Requirements for the Encoded Source Material

The requirements for the encoded source material are:

- Data must be received as an MPEG-2 transport stream (TS) over ATM Adaptation Layer type 5 (AAL5).
- The MPEG-2 TS must be ISO/IEC 13818-1 compliant.
- Program association table (PAT) and program map table (PMT) data should be repeated periodically (for example, twice per second) to allow random access.
- The MPEG-2 TS must provide a constant bit rate. The program clock reference (PCR) must be 27 MHz, plus or minus 810 Hz.
- The MPEG-2 TS must contain a single program only. The program must contain only one video packet identifier (PID), one audio PID, and one PCR PID.
- Main profile at main level decoding is supported.
- NTSC and PAL video are supported.
- Only 4:2:0 video format is supported.
- The following resolutions are supported:

  NTSC

    SIF 352x240 pixels
    HHR 352x480 pixels
    CCIR-601 720x480 pixels
    Square NTSC 640x480 pixels

PAL

SIF 352x288 pixels
HHR 352x576 pixels
CCIR-601 720x576 pixels

- Transport stream video bit rates of up to 15 Mbps are supported.
- Transport stream audio bit rates of up to 384 kbps are supported.
- Audio sample rates of 32, 44.1, and 48 kHz are supported.
- Only MPEG-1 audio is supported.

4.4 Introducing the MediaSTREAMER

The IBM MediaSTREAMER is a video server for streaming audio and video content in both analog and/or digital output formats. MediaSTREAMER solutions are designed to provide efficient storage and distribution of professional quality digitally encoded video and audio material. Utilizing the IBM MediaSTREAMER hardware and software, the solutions are designed for ease of administration, high reliability, and high quality of video.

Typically stored in MPEG-1 or MPEG-2 compressed format, audio as well as video and graphics can be distributed to local recipients using hard-wired coaxial cabling, or to more wide spread destinations using long distance cable or satellite networks. The following figure illustrates a basic MediaSTREAMER subsystem configuration.

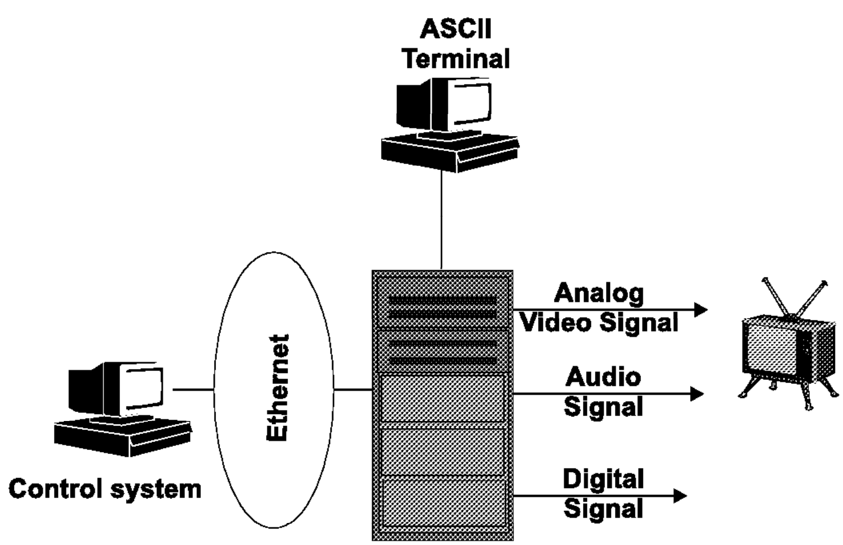

Figure 45. MediaSTREAMER System Overview

In its current form, the MediaSTREAMER is capable of sending up to 42 independent videos in NTSC (National Television Standards Committee format), PAL (Phase Alteration Line format) or through ATM in digital format up to 75 (Digital) different recipients.

The MediaSTREAMER can also be used to deliver simultaneously or independently of multiple video streams from the one source file, and even record data from satellite systems while streaming data out normally. The design of the MediaSTREAMER facilitates custom configuration for unique customer requirements.

4.4.1 MediaSTREAMER Components and How They Work

The MediaSTREAMER is comprised of two main parts: the streaming hardware and the control hardware, both with their appropriate software controls. The streaming hardware consists of hard disk, processor and video memory cache, and the control hardware consists of a control system, communications adapters and decoders.

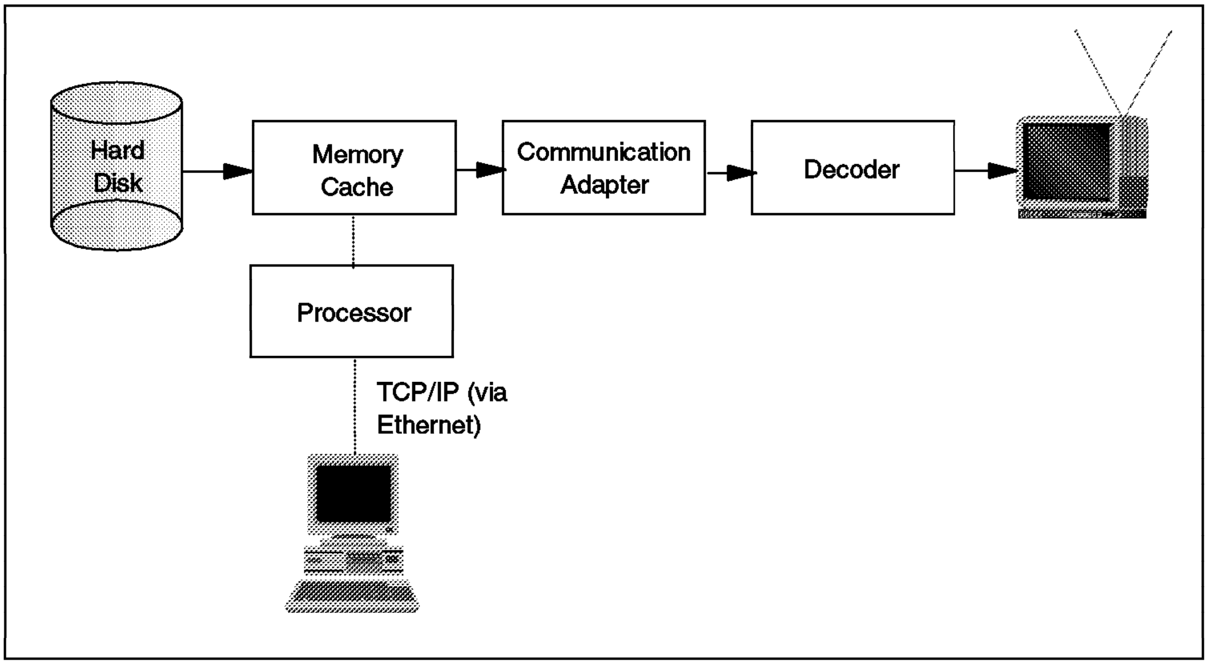

Figure 46. MediaSTREAMER System Internals

To play a stream, the Control System sends a signal to the MediaSTREAMER processor that instructs data to be pulled from disk (DASD) and sent to the cache. In the cache, the stream is mixed with other streams and sent to the controller card. This card then splits out each stream to individual decoder cards that convert the digital signal to analog.

- The Control System can be any workstation that can issue calls to the MediaSTREAMER. Calls or commands are sent in RPC (Remote Program Control) to the MediaSTREAMER connected via a TCP/IP network. When a call is issued, the MediaSTREAMER receives the call and executes the call according to system design/configuration.
- The video memory cache is random access memory (RAM).
- The processor varies depending on the desired number of streams to be delivered.
- The communications adapter sends and receives data streams. The maximum bandwidth available is 96 Mbps. The bandwidth per stream decreases as the number of streams increases.

- The decoder accepts MPEG-1 or MPEG-2 compressed signals at up to 15 Mbps, and converts to either NTSC or PAL analog signals, ready for delivery to television units.

4.4.2 System Overview

Analog video files can be encoded and loaded in a D1 format (uncompressed, broadcast quality) or in an MPEG-1 or MPEG-2 compressed mode. Encoding can be performed using an IBM Power Visualization System (PVS) or other encoding devices. The output files from the encoding process can then be transferred to the MediaSTREAMER. Data transfer can be via tape, compact disc, disk or network. The MPEG-2 video format in DVD (Digital Video Disk) is currently available in the market. With this new technology, we do not need any further compression process and we can go directly into data copy process. RPC command protocol is used to perform the loading of data on to the MediaSTREAMER. Parameters within the RPC command ensure optimal striping of the data.

Video requests from the end user can be generated from a remote PC or Set-Top Box. These requests are routed to the Control System using RPC calls.

4.4.3 Internal Architecture

The primary software component that separates the MediaSTREAMER architecture from that of the standard RS/6000 family is its advanced media file system and its ability to perform both stream I/O and file I/O on one or more ports. The following sections define some key concepts in MediaSTREAMER architecture and provides details about the file system components and their operation.

4.4.4 Stream I/O and File I/O Protocols

Stream I/O moves data in an open loop data transfer model. Many clients may receive data and there is little or no feedback during data transfer. The stream I/O protocols supported by the MediaSTREAMER are composite analog streaming and ATM AAL-5 digital streaming. The former supports only MPEG files. The latter supports any digital file but is optimized for MPEG compressed files.

File I/O moves data in a closed loop data transfer model. Each client has control over the data it receives and, in general, data checking and data retransmission are used to ensure 100% transfer integrity. Because clients can interact with the server independently of each other, broadcasting of data is not usually done.

The MediaSTREAMER supports both protocols in a variety of ways. Data may be imported (recorded/written) with either protocol and data may be exported (played/read) with either protocol. Record/Play commands support streaming, Write/Read commands support file transfer. Uniquely, the MediaSTREAMER allows these protocols to be used on the same file, even at the same time (for example, loading a file while playing it out).

4.4.5 Ports

The MediaSTREAMER streaming interfaces, whether they are physical or virtual, are called ports. Ports are selected and opened prior to using them for recording or playout. Port groups are provided to make selection a simple process. When it is not necessary to know the exact relationship between a port and a client, only the port type needs to be specified to complete a selection.

The MediaSTREAMER ports are ATM and analog. An analog port supports one stream for playout only. An ATM port is a virtual circuit which is one transfer path on the ATM adapter. As many as 75 ATM ports (virtual circuits) are provided on a single adapter in the MediaSTREAMER.

4.4.6 Assets

An asset is a file with extended properties or attributes, including compression standard (MPEG1 or MPEG-2), compression format (transport stream, program stream or system stream), resolution (SIF, CCIR601 or HHR), output format (NTSC or PAL), and bit rate (average play rate in bits/second). When an asset is staged between two MediaSTREAMERs or a MediaSTREAMER and the MediaSTREAMER Archive, the asset properties are maintained. Asset groups may be defined at configuration time to allow files to be segregated and managed independently by different applications. In this way, the MediaSTREAMER may be shared easily by different groups.

4.4.7 Application Programming Interface (API)

The MediaSTREAMER responds to commands, known collectively as the Application Programming Interface or API, from an external application or client. Session commands are used before any others to establish a context for command processing and to associate the application with a particular asset group.

Once a session is established, the client or application may issue commands to allocate ports, play videos and access other services of the MediaSTREAMER. The MediaSTREAMER uses the session context to make sure commands are appropriate and do not disturb activities of other active sessions. Applications, if they choose, may even register for session callbacks so that they don′t have to repeatedly ask for status (poll) during the execution of long commands (for example, playing a two-hour movie). In this way, pertinent information (for example, failure of an allocated port, completion of a stage operation, etc.) is provided without tying up communication resources.

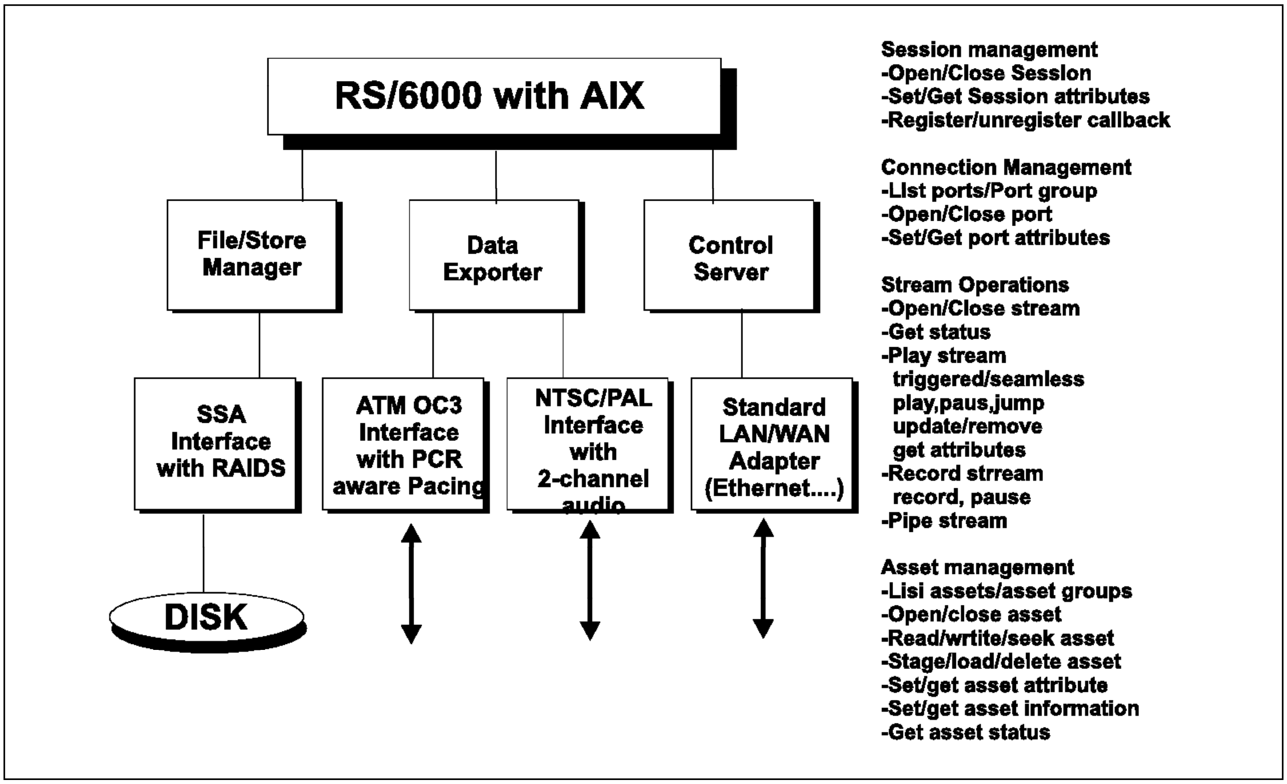

Figure 47. MediaSTREAMER Architecture and API

4.4.8 Control Server

There are a number of components that make up the MediaSTREAMER architecture. The Control Server is the component that provides the external interface, or API, to clients/applications. The API defines commands in the following groups:

- Session Management: Establishes communication with the MediaSTREAMER
- Connection Management: Determines the streaming ports available and their status
- Stream Operations: Associates ports with audio/video files and control streaming
- Asset Management: Manages files in the MediaSTREAMER (add, delete, query, etc.)

The Control Server provides a control bridge between external clients/applications and the other components in the MediaSTREAMER architecture. It manages sessions and ensures one client does not disturb the activities of another.

The Control Server determines, for the configured RS/6000 system, the available bandwidth provided by the processor and storage pool for delivering video. When clients make requests to access files or play them out, a specific amount of bandwidth may be needed to ensure smooth operation. In these cases, the Control Server allocates the bandwidth for file loading and playout and maintains it for the client. A managed bandwidth allocation of this sort is referred to as a Quality of Service (QoS) reservation.

In the MediaSTREAMER architecture, QoS reservations are required for all stream I/O operations and are optional for file I/O operations. File transfers in and out of the MediaSTREAMER may reserve bandwidth or use leftover bandwidth, that which is not reserved for QoS clients. Of course, leftover bandwidth may vary over time depending on reservations in progress. In general, QoS reservation is provided for real-time operations such as encoding files directly into the MediaSTREAMER or playing files out. In general, QoS reservation is not necessary for nonreal-time operations, such as asset staging or copying files from encoder storage.

The Control Server interprets commands on the MediaSTREAMER control interface and provides controls to the other two primary components of the architecture: the File/Storage Manager and the Data Exporter. Together, these two provide a data pump capability. This allowsfuture configurations with multiple data pumps under a single Control Server, that is, a system with a single point of control for a great number of output streams. Such a configuration takes advantage of cluster hardware or parallel processing hardware platforms, such as IBM′s SP/2 and provides simplified control for a potentially complex architecture.

4.4.9 Data Exporter

The Data Exporter is the component that isolates and manages streaming and associated ports. It sets up and provides data to both analog decoder ports and ATM virtual circuit ports. The data exporter manages output bandwidth in accordance with QoS reservation and provides an interface to the Control Server for controlling streams.

Streams come in three types: play streams, record streams and pipe streams. A play stream is one that reads one or more files and writes data to an allocated port. A record stream is one that reads data from an allocated port and writes it to a file. A pipe stream is one that reads data from an allocated port and writes it to another allocated port. The first two types of streaming require interaction with the File/Storage Manager. Pipe streams are managed by the Data Exporter alone.

Play streams have more controls than do record or pipe streams, primarily because inbound stream data cannot be controlled; it arrives at a constant rate. The commands for controlling play streams include play (immediate or triggered), pause, jump and close10. Additional play stream commands are provided to build play stream lists so that files can be concatenated together for playout. These playlist-like abilities are provided with update and remove commands. Update allows elements to be placed into the play stream and modified. Remove allows elements to be removed from the play stream.

Record streams may be used to capture inbound data into a file. Pipe streams may be used to direct inbound data to an outbound port, a feature useful for converting an inbound ATM stream to an outbound analog format. These two commands are part of an evolving playthrough architecture in the MediaSTREAMER. Future abilities will include settable pipe delays and simultaneous pipe/record capabilities.

4.4.10 File/Storage Manager

The File/Storage Manager is the component responsible for distributing data throughout the disk pool in a way that ensures continuous stream operation. The File/Storage Manager responds to asset management commands from the Control Server and coordinates data transfers with the Data Exporter.

The design of the File/Storage Manager provides the ability for a file to be opened many times and the ability to read a file while it is being written. This useful feature, also part of the evolving playthrough architecture, provides low latency data staging: the ability to satisfy request(s) to view a file while it is being copied into the MediaSTREAMER. Such capability is extremely valuable in configurations where MediaSTREAMER Archives provide data to MediaSTREAMERs, where centralized MediaSTREAMERs provide data to local MediaSTREAMERs or where MediaSTREAMERs are managed to copy data from each other.

The File/Storage Manager must manage data distribution in a way that ensures the bandwidth capabilities of the MediaSTREAMER configuration are met. Therefore, when storage is added to (or removed from) the MediaSTREAMER, the File/Storage Manager must redistribute data accordingly. In essence, all the data in the storage pool must be read and rewritten so it is desirable to minimize stored content at the time of reconfiguration. In a typical situation, where the capacity on a 200 GB machine is doubled when it is half full, about 7 hours are needed to reorganize the data.

4.4.11 Storage Pool

The MediaSTREAMERs SSA disk attachment technology provides a large capacity (over a terabyte) storage pool for video storage. The large number of disks in such a configuration introduces the need for RAID protection, since disks are active mechanical components which are more prone to failure than other components in the subsystem. The RS/6000 disk drawer design provides protection from a number of disk related failures with: RAID, a hot spare disk, a redundant power supply, a redundant cooling fan and hot replacement of a failed disk. These protective features result in a MediaSTREAMER that is highly reliable even with large pools of disks.

The MediaSTREAMER configures disks in groups of five, each group of five making up a RAID array. Three of these arrays can be configured in one disk drawer, which holds 16 disks, allowing room for a hot spare. The number of disk arrays in a configuration will depend on the need for storage capacity, streaming capacity and tolerance for video disruption. The disk pool may be ordered in either a temporarily degrading or fully nondegrading configuration.

In a temporarily degrading configuration, a disk failure will cause a loss of system bandwidth that may disrupt video output, depending on the system load at the time of failure. The loss of bandwidth lasts only as long as it takes to reconstruct data on the failed disk on one of the internal hot spare disks.

In a fully nondegrading configuration, which has more storage bandwidth, a disk failure causes no disruption in video output. Replacement of the failed disk is nondistruptive in both configurations.

The number of RAID arrays required for temporarily degrading and fully nondegrading configurations depends on the bandwidth required of the

MediaSTREAMER system. For a maximum bandwidth configuration, seven RAID arrays provide fully nondegrading capability and two RAID array provide temporarily degrading capability.

4.4.12 MediaSTREAMER Type and Model

IBM MediaSTREAMER Analog System-A11: A smaller video server configuration, has two standard analog NTSC/PAL video output streams (expandable to 42 streams) and 18 GB of content storage. A maximum of three additional drawers can be added in the rack containing the processor for expansion: up to three content storage drawers (maximum of 216 GB), two analog systems (up to 42 streams), or a mixture of content storage and analog subsystems. The A11 has two Microchannel adapter (MCA) slots available for expansion One ATM-155 Video Streaming Adapter can be added to support up to 75 digital streams (requires two MCA slots). Each analog subsystem require one MCA slot.

IBM MediaSTREAMER Analog System-A12: A larger video server configuration that has eight standard analog NTSC/PAL video output streams (expandable up to 42 streams) and 54 GB of content storage (expandable to 216 GB maximum). A maximum of three drawers can be added for expansion in the rack containing the processor: up to three content storages (216 GB), two analog subsystems (up to 42 streams), or a mixture of content storage and analog subsystems. The A12 has five MCA slots available for expansion. One MCA slot is reserved for a network adapter. Two ATM-155 Video Streaming Adapters can be added to support additional connectivity (requires two MCA slots each).

IBM MediaSTREAMER Digital System-D11: A smaller video server configuration, has one ATM-155 Video Streaming Adapter (up to 75 digital streams) and 18 GB of content storage. A maximum of five drawers can be added for expansion in the rack containing the processor: up to five video content storage drawers (maximum of 270 GB), one analog subsystem (up to 14 analog streams), or a mixture of content storage and analog subsystems. The D11 has one MCA slot available for expansion.

IBM MediaSTREAMER Digital System-D12: A larger video server configuration that has one ATM-155 Video Streaming Adapter (up to 75 digital streams) and 54 GB of content storage (expandable to 216 GB maximum). A maximum of four drawers can be added for expansion in the rack containing the processor: up to four video content storage drawers (maximum of 270 GB), two analog subsystems (up to 28 analog streams), or a mixture of content storage and analog subsystems. The D12 has four MCA slots available for expansion. One MCA slot is reserved for a Network Adapter. One additional ATM-155 Video Streaming Adapter can be added to support additional connectivity (requires two MCA slots each).

IBM MediaSTREAMER Content Storage Expansion C10: A separate R00 rack is available to expand the content storage for A11, A12, D11 and D12 systems. Six content storage drawers (maximum of 324 GB) can be installed in each expansion rack. Each drawer contains 54 GB of content storage. Up to three content storage packages can be attached to a MediaSTREAMER System solution. One available MCA slot in the A11, A12, D11 and D12 is required to install a SSA 4-Port RAID Adapter that must be ordered for each C10.

MediaSTREAMER capacity and throughput of the analog and digital solution can be increased by adding content storage drawers in the processor rack or in an expansion rack, an additional analog subsystem drawer in the processor rack, and an additional ATM-155 Video Streaming Adapters. Each analog subsystem drawer can have from 2 to 14 MediaSTREAMER Audio/Video Decoder cards.

4.4.13 The IBM Multimedia Server Solution

Much interest is currently being shown in applications that fall into the general category of digital library. A *digital library* is a machine readable collection of materials, such as those that would be found in a traditional library, together with organizational information necessary for users to find specific information. Unlike the traditional library however, a digital library contains multimedia data, such as images, sound and full-motion video, in addition to text data. Applications of this technology include in-store advertising/shopping, education, kiosks, home shopping, online libraries, and more. One major common component of all these applications concerns the storage and retrieval of large amounts of heterogeneous data objects. This component can be addressed using a solution built around the MediaSTREAMER that is called the *IBM Multimedia Server Solution Offering*. Other common components, such as messaging and transaction processing, are addressed by other IBM products.

In order to meet the diverse digital library requirements, an architecture for the IBM Multimedia Server Solution Offering has been specifically developed. The key element of this architecture is called a *multimedia node*. Each multimedia node contains four components:

1. A *library server*
2. One or more nonvideo *object servers*
3. One or more *stream (video) servers*
4. One or more storage subsystems for the above

Each of the components is separately and independently scalable allowing each node to be individually configured.

Objects managed in the multimedia node are indexed in the library servers, and are delivered from object servers. A library server is analogous to a greatly enhanced card index in a library; it knows many things about the object, including where it is located. Object servers are analogous to bookshelves; they physically hold objects. A stream server is an important special case of an object server, optimized to cost-effectively play back real-time full-motion video streams.

A single, central node is designated as a master site. This holds the unique, master data index. (If other distributed nodes have local library servers, these indices are volatile and contain simple subsets of the master index.)

If the master site also contains a single master copy of each data item, this allows us to treat other sites as simple blob caches and the data at the sites also becomes volatile. If the local cache fills up, or data objects ″time out″, then they can simply be deleted from the distributed node.

The IBM Multimedia Server Solution Offering is a set of standard IBM products packaged as a service offering designed to address the data storage and manipulation requirements encountered in a wide range of digital library and Video-on-Demand applications.

The IBM Multimedia Server Solution Offering is essentially a distributed client/server model that allows servers to be organized in a flexible hierarchy, and implements mechanisms that allow centralized control to simplify the problems associated with the backup, recovery and management of large distributed systems.

4.4.14 Characteristic of Digital Library System

All digital library type applications have a number of characteristics and requirements that are common:

- Large numbers of data objects to be stored
- Arbitrary mix of data types
- High storage and bandwidth requirements

Typically, the asset value of the content (information) stored in a digital library is high; in many cases, it is, in fact, a key corporate asset (this is true, for example, of a film company). Usually, revenue is earned by allowing users to view parts of this content for a fee (in our example, pay-per-view movies, for instance).

Typically, the content is copyrighted in some way, and therefore viewing or movement of the data may incur royalty charges. Thus, data movement needs to be carefully managed and recorded. Workflow management techniques provide a convenient mechanism for many of the routine data entry, deletion or movement functions.

For royalties to be paid, information on data usage must be collected, and therefore, typically, there are requirements for capturing this information. This may be conveniently handled by a transaction processing system (for example, CICS) with data forwarding to a central processing site being provided by a message forwarding system (such as IBM′s Manufacturing Quality System).

There may be additional requirements for security involved if the data is sensitive (that is, personal or medical data) or if financial information, such as credit card numbers, is involved.

4.4.15 Hardware and Software

The MediaSTREAMER is a flexible hardware/software package containing an RS/6000 uni-processor with MediaSTREAMER for AIX pre-loaded software. The uni-processor may be either a 39H (adapter slots = 4) or an R20 (adapter slots = 8) packaged in a 19″ rack with storage and I/O adapters.

RS/6000 with AIX

- 39H with 4 adapter slots (and a built in Ethernet Adapter)
- R20 with 8 adapter slots
- Analog throughput: 100 Mbps
- Digital throughput: 120 Mbps
- I/O ratio: 33 % inbound data max for a fully loaded system
- Video Control command execution within 250 ms

- Initial title per-roll (cue) delay < 3 seconds
- Minimum video clips length: 10 - 20 seconds

Streaming Adapters

- 1 - 42 analog adapters with NTSC/PAL output
- 1 - 2 ATM streaming adapters with PCR aware AAL-5 protocol

TCP/IP Adapters for File I/O and Application control

- ATM-155 Turboways adapter (155 Mbps OC-3c multimode fiber or copper)
- ATM-155 Video Streaming adapter (155 Mbps)
- Ethernet 10Base-T
- Ethernet /FDX adapter (10 Mbps)
- Ethernet High Performance adapter (10 Mbps)

Storage

A single SSA RAID5 adapter supports up to 96 disk drives = 6 disk drawers. The configuration of the disks is in groups of five, each five being a RAID5 array. Fully configured, MediaSTREAMER 1 supports up to fives drawers in three expansion racks for a total of 23 drawers = 1.242 terabytes of data storage. Attributes are as follows:

- 1-3 RAID arrays per drawer:
  - 18 GB storage capacity per array
  - 54 GB storage capacity per array
- Optional hot standby disk per drawer
- Redundant power supply per drawer
- Redundant cooling fan per drawer
- 6 drawers per expansion cabinet (rack)

Other Hardware required

- Network attached graphical workstation (or PC with a Web browser) for configuration and administration:
  - Optional system monitoring with NetView/6000 (via SNMP)
  - Optional performance monitoring with IBM Performance Toolbox for AIX
- Serially attached ASCII terminal for service and setup

Clients/Receivers

- ATM AAL-5 streaming

  IBM 8260 Nways Multiprotocol Switching Hub or 8285 Nways ATM Workgroup Switch with 8-MPEG2 Decoders ″Video Distribution Module″ Blade
- NTSC/PAL analog:
  - Standard NTSC/PAL Television Monitor
  - Cable Head-End equipment utilizing NTSC/PAL composite inputs
  - PC-TV clients

4.4.16 Network Components

The MediaSTREAMER is a storage device designed to provide online media storage for delivery to a network. When the MediaSTREAMER and control system are connected via a network, collectively they form a media delivery network. To connect the MediaSTREAMER and control system, the following network components are required:

- Communications adapters
- Cabling
- Communications protocol
- Communications access units
- Connectors, terminators and transceivers

4.4.17 Example Environments

There are two example environments in which MediaSTREAMER operates: Network and Broadcast Automation Environment

- **Network Environment:**

  The Network illustration shows the MediaSTREAMER network environment as it:

  – Interacts with one or more video creation systems, using FTP or content management application to receive encoded multimedia assets, such as video files.

  – Interacts with management programs to add, delete, and schedule multimedia assets.

  – Requests that assets stored on the MediaSTREAMER archive across the network be staged for transmission to clients at a pre-scheduled time.

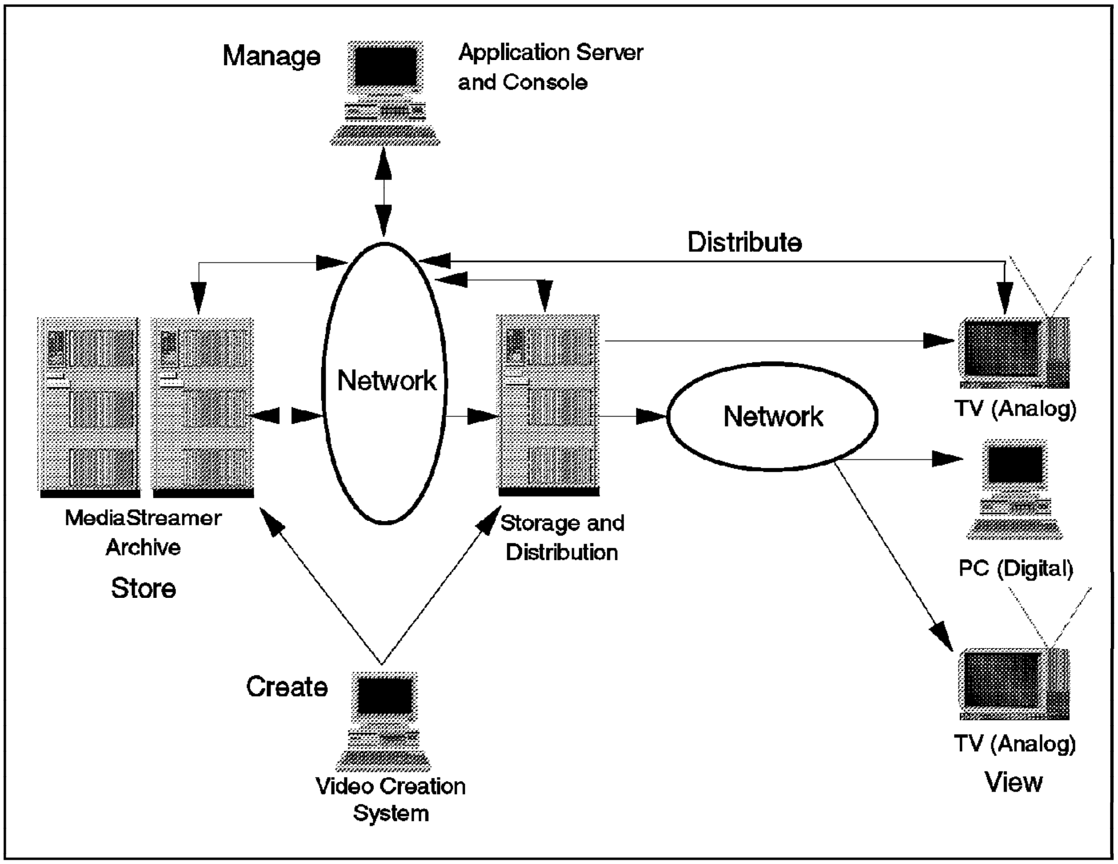

Figure 48. Network Environment

The MediaSTREAMER Archive is used to store large quantities of multimedia content. The content is translated to digital format by an encoder and stored on tape cartridges in the tape library. The digital assets are staged (copied) from the MediaSTREAMER archive to the storage and distribution unit when it is time for them to be played.

4.4.17.1 Broadcast Automation Environment

The Broadcast Automation Environment (figure 4) illustration shows the MediaSTREAMER environment using a broadcast automation system.

For the broadcast industry, the use of digital video server provides the advantages of digital fidelity and computer-based automation. Analog video tapes, and the need to move them around with ″sneakernet″, are replaced by disk storage and data file transmission. These general benefits, in video servers such as the MediaSTREAMER, enhance a variety of applications.

The adherence of open system standards allows the MediaSTREAMER to be a building block in a number of video solutions with complementary products from IBM and other vendors. IBM MediaSTREAMER Archive provides automated storage and retrieval of large amounts of digital data in a robotic tape system. IBM′s Master Automation Bridge allows popular automation systems such as Louth System to control both MediaSTREAMER and MediaSTREAMER Archive.

IBM′s Digital Library provides a way to manage (locate, search and protect) digital data objects worldwide, including text, image and audio or video objects.

IBM 8300 Video Access Node (VaN) and IBM Video Distribution Module (VDM) products enable high quality video conferencing or distance learning over local or wide area ATM networks. IBM 8260 Nways Multiprotocol switching hub with the new Video Distribution Module (VDM) is used to interoperate with any other switching devices. When any or all are coupled with the MediaSTREAMER, powerful solutions may be built for broadcast and other industries:

- Traditional broadcast (analog tape replacement, ad insertion, playout automation)
- Corporate intranet broadcast (education/training, conferencing, information services)
- Home digital services (education/entertainment, conferencing, information services)
- Home digital services (education/entertainment, home shopping, stock market)
- Video capture and search (security cameras, conferencing, news on demand)
- Special applications (kiosks, karaoke and hospitality, hospitals, prisons, casinos).

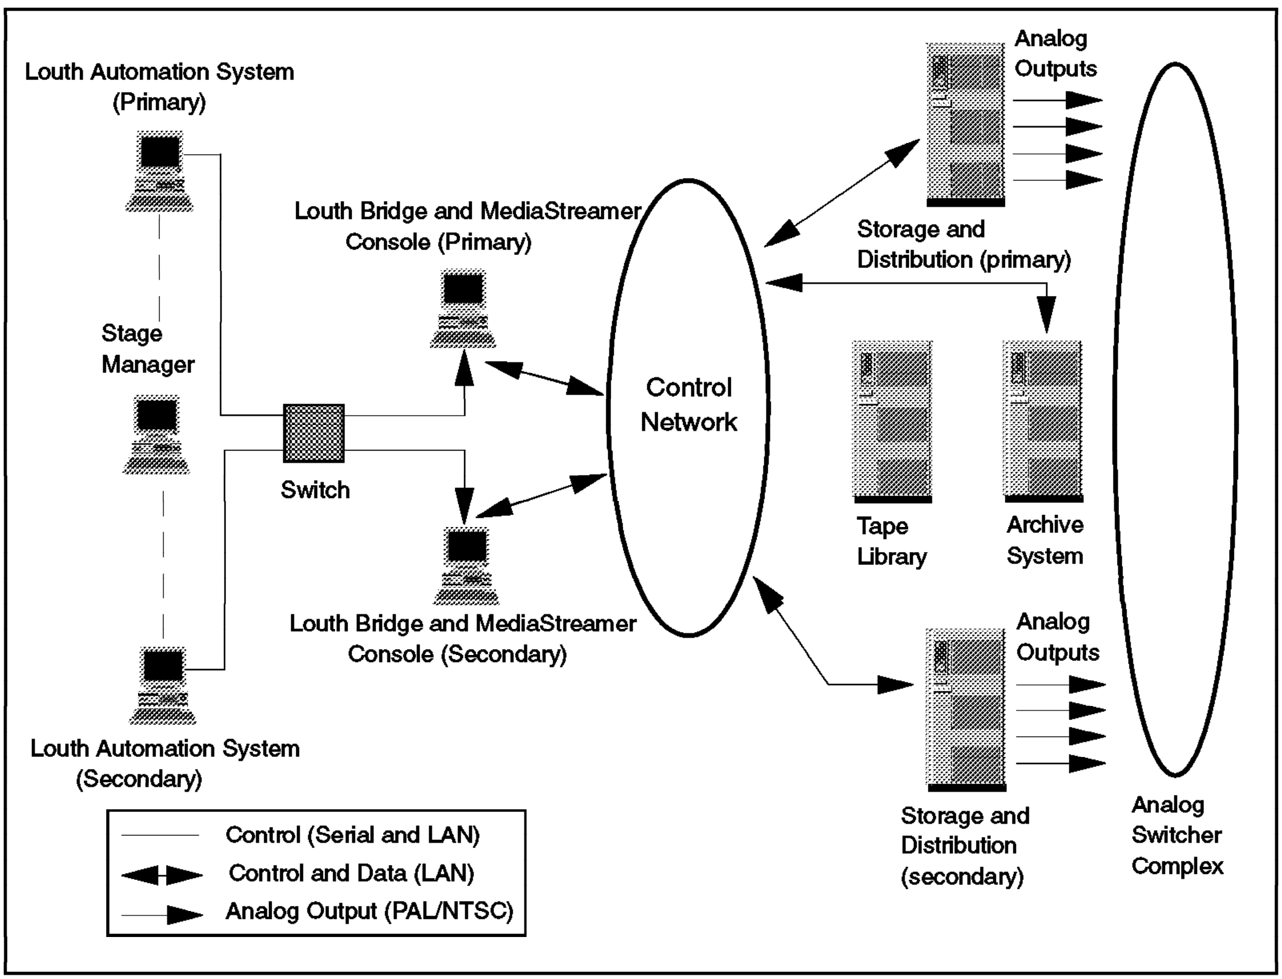

Figure 49. Broadcast Automation Environment

There are some broadcasting automation systems available on the Market (Alamar, Louth system, etc.) and they are all can be connected to the MediaSTREAMER as it is described in the following illustration.

The Broadcasting system sends commands, such as play and stop, through a serial interface to the broadcasting bridge. The broadcasting bridge then accepts and interprets these commands, issues the appropriate commands to the MediaSTREAMER or to the MediaSTREAMER archive, and returns any required status information to e the controller. The Stage Manager monitors the play list (containing a list of video assets, for example) and issues commands to copy a multimedia asset from the MediaSTREAMER archive to the MediaSTREAMER. The Stage Manager can also delete multimedia content from the storage and distribution unit if the content is not immediately needed.

The other important feature is the capability to handle *frame accuracy* of video clips. The MediaSTREAMER is already equipped with this function to ease the changes between one video to another within a sub-second period of time. Broadcasting companies need this feature to maintain the smoothness, easiness and accuracy of video clips exchange.

4.4.18 Capacity Planning for MediaSTREAMER

The MediaSTREAMER is configured to the customers′ requirements. Items that define the MediaSTREAMER configuration are:

- ATM
- Video type (NTSC or PAL)
- Network (Ethernet)
- Number of streams required (number of television units)
- RAID5 (optional)
- Amount of video to be stored (disk space required)

When calculating the amount of disk space required on the MediaSTREAMER to hold video data, it maybe helpful to use the following as a guide. Approximately one minute of video(and audio) compresses to 45 MB of data. Of course, this will vary depending on the amount of audio and frame detail in the video.

4.4.19 Capacity Planning for Control System

The control system is a workstation the customer provides that requires the following hardware and software.

The control system will need to store the following software:

- Operating system
- TCP/IP
- Control Interface
- Any other desired application software

The control system workstation will comprise the following hardware components:

- 486 or higher processor
- 12 MB RAM minimum
- Mouse
- Keyboard

- Display (IBM 8512 or above)
- Communications adapter (ATM or Ethernet)
- Communications cable
- Hard disk (540 MB)

It is suggested that the control system have enough hard disk space to hold the following software:

Table 5. Hard Disk Requirements

Software	**Example**	**Disk Space Required**
Operating System	AIX Version 4 or OS/2	30-70 MB
Communications	TCP/IP for AIX or OS/2	8 MB (base)
Communications	TCP/IP Programmers Toolkit	Included
Communications	Network file system (NFS)	Included
Control Interface	API	1 MB
Sample Code (see Note 2)	Examples	1 MB

Note: Hard disk space is used temporarily to transfer data to the MediaSTREAMER. This is one way of transferring video to the MediaSTREAMER.

Depending on how the video is transferred to the MediaSTREAMER, it may be necessary to have enough disk space on the control system to download a video (digital file) before it is transferred from an encoder to the MediaSTREAMER. Video can be transferred in other ways, such as using tape, across a network or directly from a CD-I disc.

4.4.20 Other Considerations

The following topics should be considered when planning for MediaSTREAMER use and its operating environment. All environmental issues are the customer′s responsibility. IBM support may be needed as assistance in planning a suitable working environment for the MediaSTREAMER.

- **Encoding**

  To use the MediaSTREAMER as a media server, video data must be stored on the MediaSTREAMER in a useable format. MPEG-1 and uncompressed video formats are supported. If the data needs to be compressed, such as compressing video into MPEG-1 format, this is the customer′s responsibility. If the video is already compressed, such as video stored on a compact disc in CD-I format, the data can be read from the CD-I disc to the MediaSTREAMER. No further compression would be required, as the data is already in MPEG-1 format. This is also applied for DVD with MPEG-2 format.

  The IBM 7245 Power Visualization System (PVS) can be used to compress video in to MPEG-1 format. The PVS is also capable of compressing and editing video and audio in real time.

  Once compressed, the digital video file needs to be transferred to the MediaSTREAMER. This can be done from a tape unit or hard disk on the control system performing a file transfer via FTP command. As the file is in digital format, it may also be sent over networks.

- **RAID5**

Video data is striped over many disk drives, and it is highly recommended that RAID5 (Redundant Array of Independent Disks) storage technology be used to protect data in the event of a disk malfunction. In the MediaSTREAMER, data is stored across many disks. If the data on a disk is no longer able to be used, then all the videos that had part of them stored on that disk will be lost. Instead of losing a disk with one video on it, many videos will be affected. RAID5 technology is an available option when ordering the MediaSTREAMER.

- **Storage**

  The MediaSTREAMER is capable of storing up to 216 GB of data (maximum).

- **Cable**

  1. It is necessary to have planned the BNC to deliver PAL/NTSC analog signals, the ATM fiber optic cable to deliver digital signals, and the Audio cable.

- **Radio Emissions**

  The MediaSTREAMER radiates radio frequency energy. It is important that environmental factors be considered during planning and installation. For example, if the MediaSTREAMER were to be placed in an airplane, it may interfere with the transmission of other radio signals. To contain the radio emission, the MediaSTREAMER must be installed and used in accordance with the instruction manual. If required, the MediaSTREAMER can be housed in a shielded environment to contain the radio emissions.

Chapter 5. Configuring the VaN and the VDM

We describe in this chapter the configuration steps required for the Video Access Node and the Video Distribution Module. Several configurations scenarios are also included to illustrate interoperability between the VaN, VDM and the MediaSTREAMER.

5.1 Basic VaN Setup

This section outlines the configuration process of an IBM 8300 Video Access Node (VaN).

The following parameters must be planned in advance and must be determined before starting the configuration process.

- **ATM Address**

  You must assign the full 20-byte ATM address for the VaN box, following the ATM addressing standard, as shown in Figure 50.

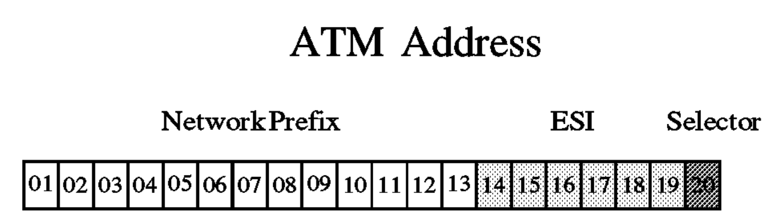

Figure 50. ATM Address Format

  – Network Prefix:

    This 13-byte parameter must be the same as defined in the ATM switch of your network.

  – End System Identifier (ESI):

    The burned-in MAC address of the ATM adapter or a locally administered address (LAA) will be used for this 6-byte parameter.

  – Selector byte:

    This byte will be defined later, when you configure the logical terminals of your VaN. The VaN uses a value of 0 for the selector byte for its TCP/IP over ATM support. Each one of the logical terminals that you create for your VaN will have a different selector byte.

- **IP Host Name and IP Address**

  The Video Access Node uses TCP/IP for the control protocol. So, you also have to define an IP address and host name for the ATM interface. The VaN also comes with a built-in Ethernet adapter. This is a useful feature if you want to connect the boxes to your local area network, either for management purposes or any other remote application. If you decide to use the Ethernet (or an extra token-ring) adapter, then you must choose IP addresses and host names for these interfaces also.

Note that the IP parameters of the local area network adapter are different from the IP parameters of your ATM adapter.

- **ARP Server**

  Since you will be using TCP/IP over ATM, you need one ATM ARP server to resolve the ATM addresses to IP addresses in your ATM network. You should use one of the VaNs as the ARP server and all the other VaNs and H.310-compliant devices in the network as ARP clients.

- **Devices and Logical Terminals**

  Encoders and decoders on a VaN are named and grouped into logical terminals by records in a file called AVS.CFG, which by default resides in the \ibmavs\etc directory.

  The logical terminals for your VaN can be configured in many ways depending on the combination of encoder and decoder adapters you ordered. Logical terminals can be of the following types:

  – ROT (Receive-Only Terminal): Consists of one decoder.

  – SOT (Send-Only Terminal): Consists of one encoder.

  – RAST (Receive-And-Send Terminal): Are made of both an encoder and a decoder.

  For the task of creating logical terminals, you must determine how many different configurations you want to make available on your network. For example, if your VaN has one encoder and one decoder, you have three possible logical terminals you can define:

  – One SOT

  – One ROT

  – One RAST

  The SOT and the ROT can be active at the same time, but if either is active, the RAST cannot be activated. The AVS software prevents the activation of a terminal when one of its devices is busy. You might want to configure your VaN with all three terminal definitions so that you can either use it as a videoconferencing terminal (RAST), or you can use the encoder and decoder independently.

  Each logical terminal uses an assigned selector byte value to generate its own ATM address from the VaN base ATM address. These selector values can be any number from 1 to 127. Each logical terminal selector must be unique within the VaN.

5.1.1 Configuring the ATM Address of the VaN

The following steps describe the procedures to configure the parameters discussed in the previous section:

1 Open the OS/2 System folder at the initial screen/desktop (see Figure 51 on page 97).

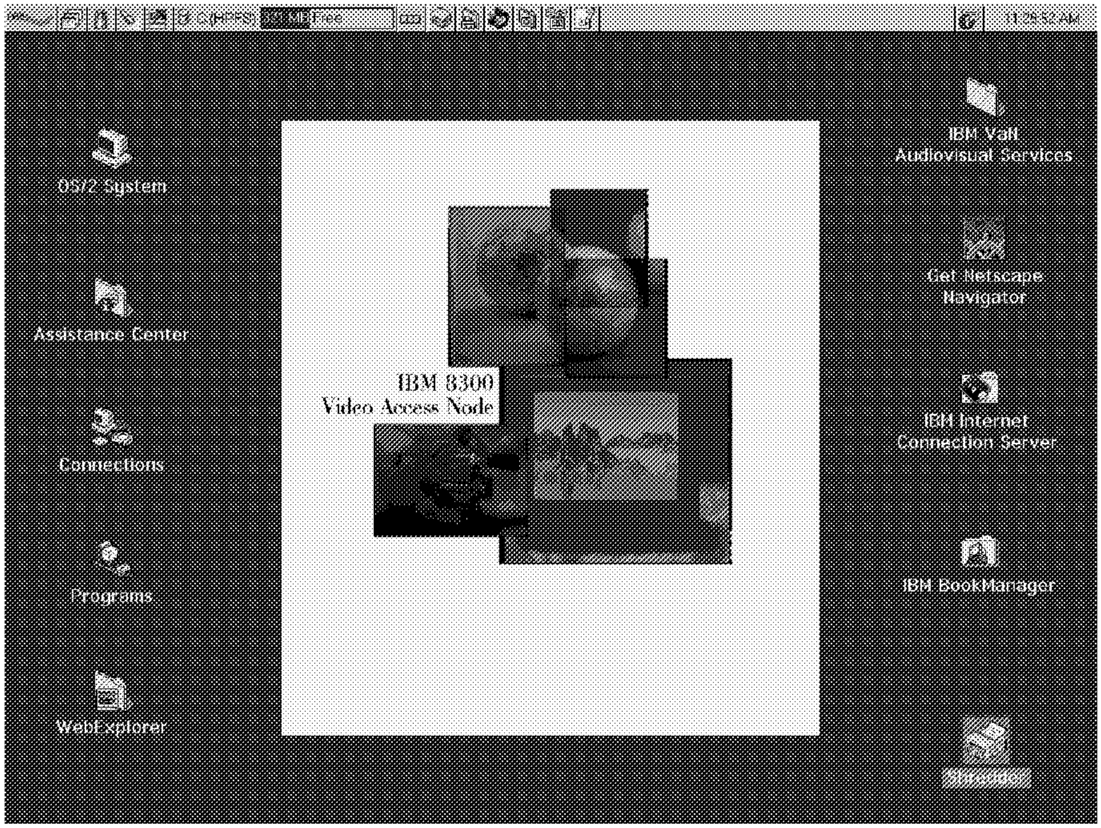

Figure 51. Initial Screen

2 The OS/2 System panel appears (as shown in Figure 52). Double-click on the **System Setup** icon.

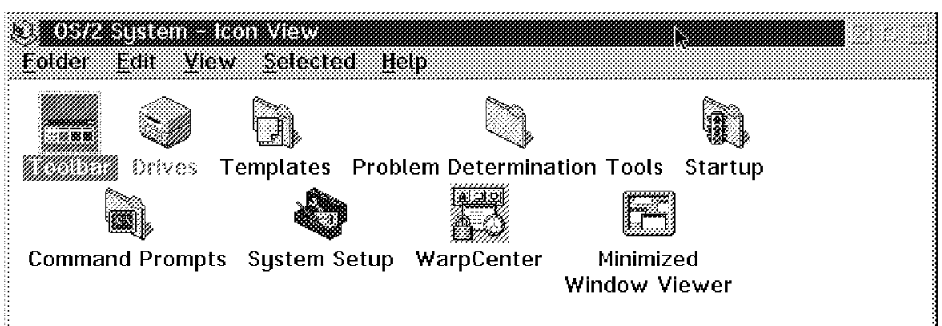

Figure 52. OS/2 System Folder

3 The System Setup panel appears, displaying the icons for the installed software (as shown in Figure 53 on page 98).

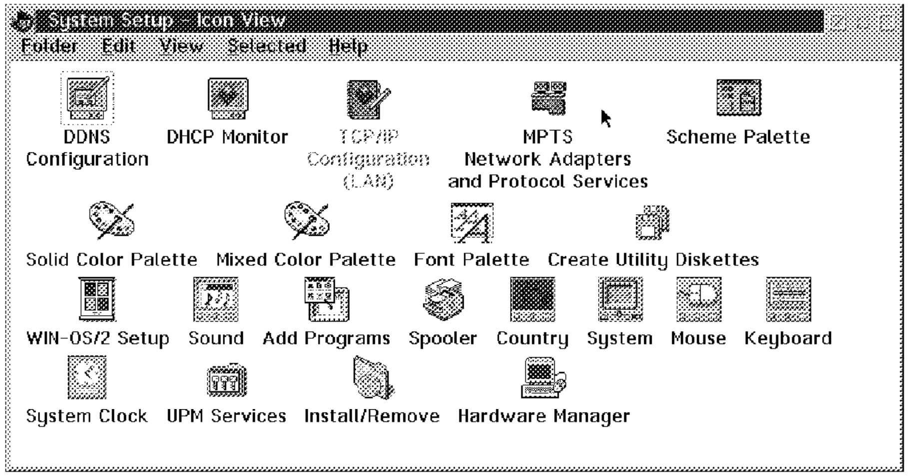

Figure 53. System Setup Folder

4 Start MPTS by double-clicking on the **MPTS** icon. If the MPTS icon is not available, execute it from the command line, by typing `C:\IBMCOM\MPTS`. The Multi-Protocol Transport Services logo appears (as shown in Figure 54).

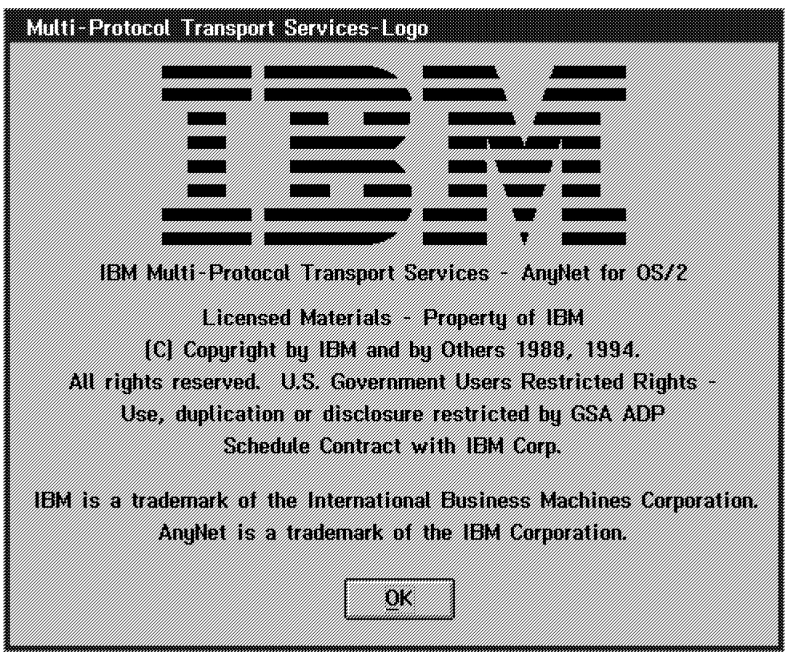

Figure 54. Multi-Protocol Transport Services Logo

5 Click on **OK**. The Multi-Protocol Transport Services panel (as shown in Figure 55 on page 99) is displayed.

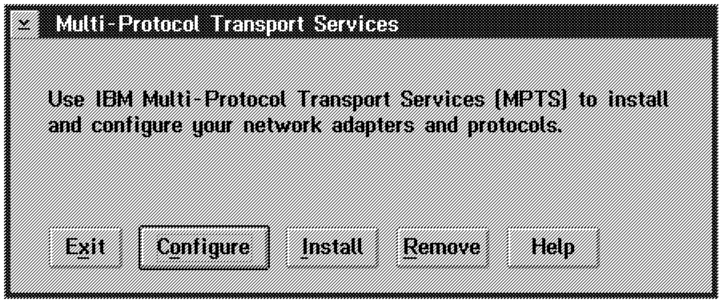

Figure 55. Multi-Protocol Transport Services Panel

6 Select **Configure**. The Configure panel (as shown in Figure 56) is displayed.

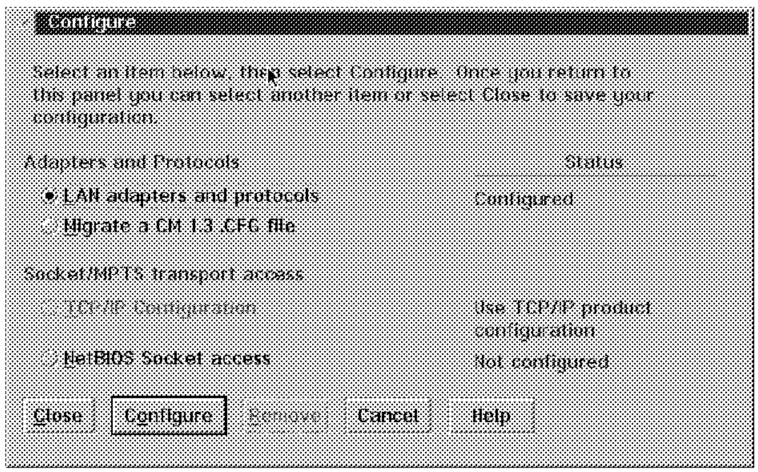

Figure 56. LAPS Configure Panel

7 Ensure that the LAN adapter and protocols radio button is selected and click on **Configure**.

The following informational message is displayed:

```
LAN adapter and Protocol Support
is reading Configuration information.  Please wait.
```

This is followed by the Adapter and Protocol Configuration panel as shown in Figure 57 on page 100.

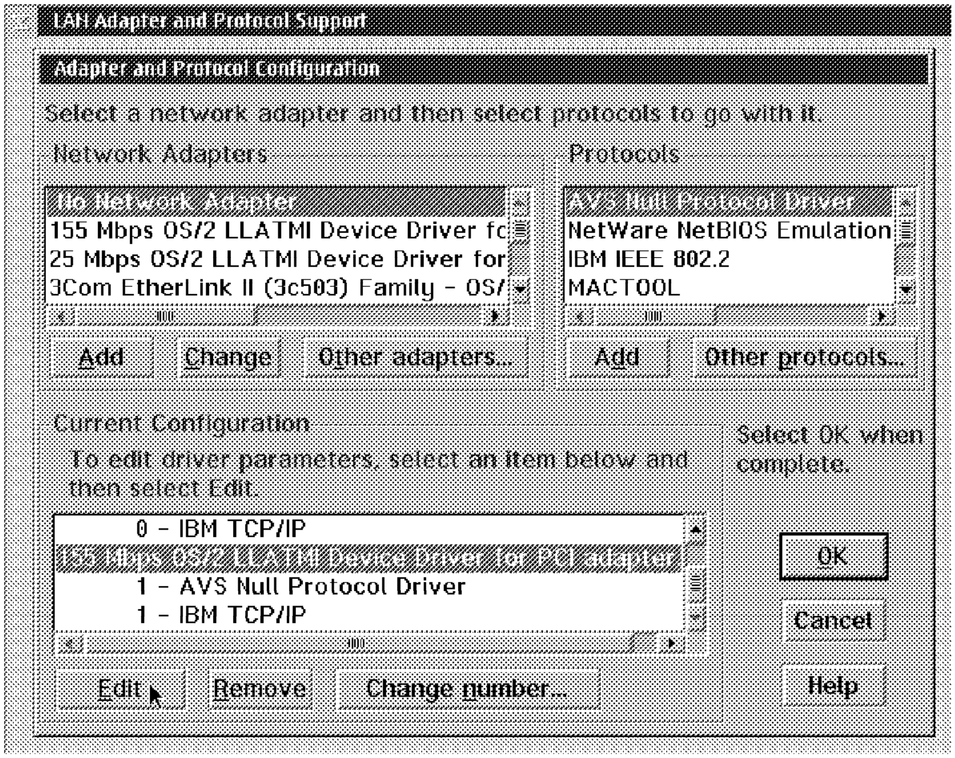

Figure 57. Adapter and Protocol Configuration Panel

8 From the Current Configuration group box, ensure the ATM adapter is highlighted and select **Edit**. The Parameters panel for the installed ATM adapter will be displayed. Figure 58 on page 101 shows the panel for the 25 Mbps ATM adapter and Figure 59 on page 101 shows the panel for the 155 Mbps ATM adapter. These panels display all the driver parameters that you can edit.

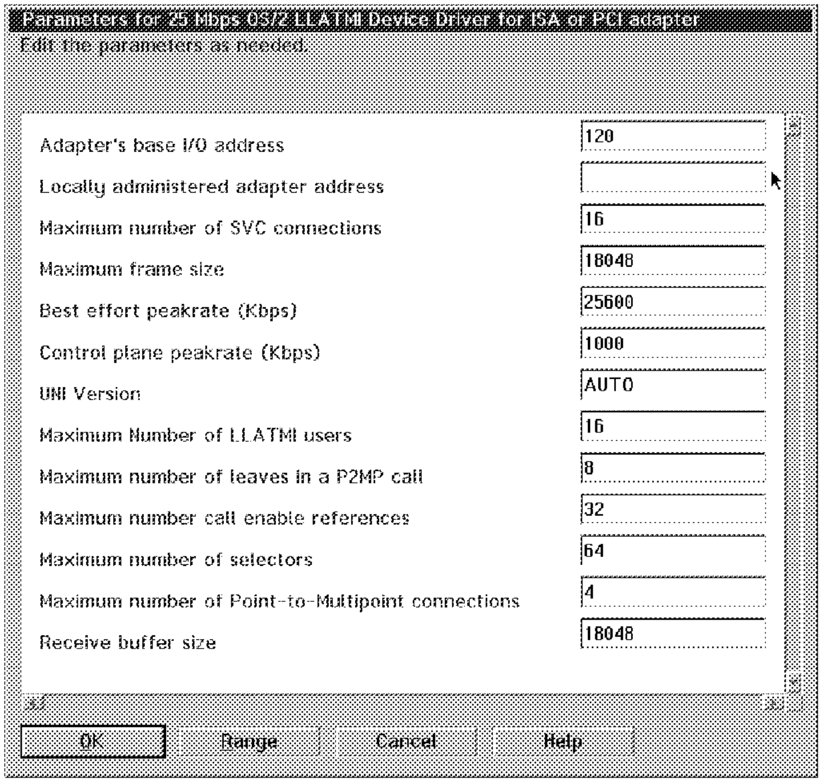

Figure 58. 25 Mbps ATM Adapter Parameters

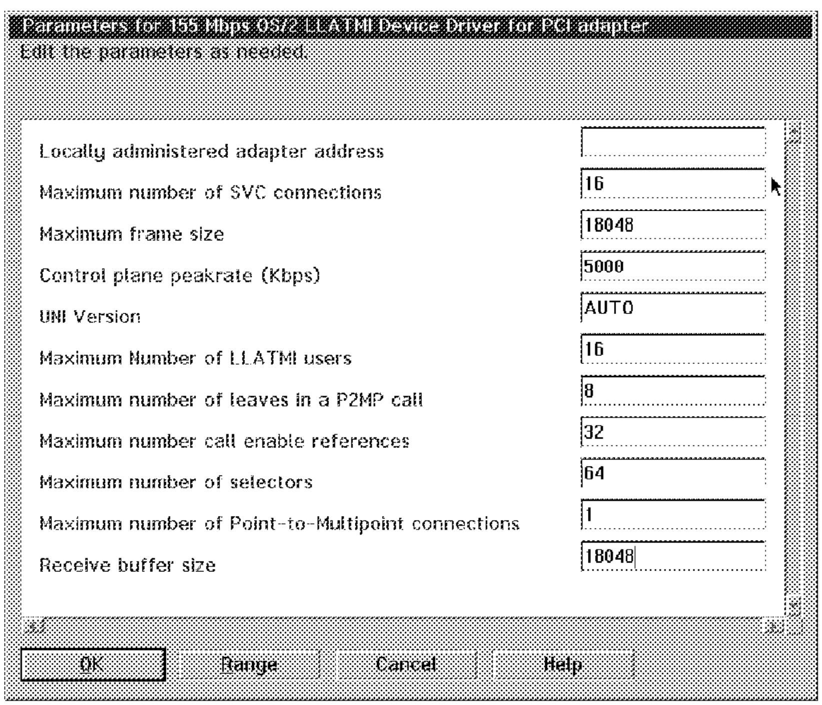

Figure 59. 155 Mbps ATM Adapter Parameters

9 The Locally administered adapter address corresponds to the ESI field of the ATM address. If you want to use a locally administered address instead of the universally administered address (burned-in MAC address), type the 12-character, hexadecimal address that you defined. If you want to use the universally administered address, leave this address blank.

10 The remaining parameters have default values assigned. If you want to change the default parameters for the adapter, type the appropriate parameter values.

> **Note**
>
> The defaults are set to maximize performance for most installations; changing them is not required. Change the defaults only if you know that you have a specific need to alter their values.

11 When you have finished editing the driver parameters, select **OK**. The Configure Workstation panel reappears.

12 Select **OK** to return to the Multi-Protocol Transport Services panel, and then select **Exit**. The update CONFIG.SYS panel appears.

13 Make sure that the disk drive shown on the panel is the same as the disk drive that contains the CONFIG.SYS file.

> **Note**
>
> The message `Exiting MPTS - Please wait` may take a considerable amount of time. Do not power off your workstation until advised to do so.

Select **Continue**. When the LAPS program has finished updating the CONFIG.SYS file, select **OK**. The Exiting MPTS panel appears.

14 Select **Exit** to exit the MPTS program.

15 Shut down your OS/2 system and restart (reboot) your computer to complete the driver installation. Do not restart the workstation until you are advised that it is safe to do so.

5.1.2 The AVS.CFG File

The C:\IBMAVS\ETC\AVS.CFG contains the device and logical terminal definitions for the VaN. The AVS will read this configuration file when it is started on the VaN.

The AVS.CFG file is shipped with each VaN and matches the number of encoders and decoders installed. You probably want to edit it to change the terminal names or to include modifications on the number of codec cards installed at a later time.
The AVS configuration file contains records that define:

- H.245 operational parameters
- The base value used in combination with ATM selector values to determine TCP port numbers for logical terminals
- Whether the VaN may respond to inbound requests without application-level control
- The audiovisual encoding and decoding devices available to the VaN
- The logical terminal configurations to present to the ATM network

Commands for performing these tasks are summarized in Table 6 on page 103 and are described in more detail in the next section.

Table 6. AVS Configuration File Summary

Record	Function
device	Describes an audiovisual encoding or decoding device
h245maxsessions	Defines the maximum number of concurrent H.245 control sessions
h245time-out	Defines the time-out interval for each H.245 control session
lterm	Defines a logical terminal
supervised	Indicates that the VaN will only allow inbound requests when approved by a local application program.
tcpportbase	Defines the base value for determining TCP port numbers
unsupervised	Indicates that the VaN will honor inbound requests regardless of whether an application program is registered.
#	Comment

5.1.2.1 The device Record - Audiovisual Encoding or Decoding Device

This record describes a real or virtual encoding or decoding device.

Format

```
                                                           ┌─MNTSC358────┐
►►──device──devname──┬─┬─encoder─┬──┬─cardNumber──┬──────────┬─┼─────────────┼─┬─┬──►◄
                     │ └─decoder─┘  │             └─backup─┘ └─videoFormat─┘ │ │
                     │              └─virtual──┬──────────────────────────┬──┘ │
                     │                         └─exename──┬─────────────┬─┘    │
                     │                                    └─parmstring─┘      │
                     │                     ┌─D:\$TEMP$.M2T─┐                   │
                     └─┬─fileplayer────┬───┼───────────────┼───────────────────┘
                       └─file recorder─┘   └─mpegfilename──┘
```

Parameters

devname

The name by which this device will be referenced in the VaN. This name is case-sensitive and must be between 1 and 63 characters long. Any characters are allowed. If blanks, hyphens, or any OS/2 delimiters are used, the name should be enclosed in double quotes. For example, all of the following are valid device names:

```
classroom5decoder
Classroom #5 - Left side of room
```

encoder

Indicates that the device is an encoder (generates transport streams).

decoder

Indicates that the device is a decoder (receives transport streams).

cardnumber

A single-digit number that correlates to the discovery order of the physical adapters. Each type of adapter (encoder and decoder) is "discovered" by the device drivers during boot. The order of discovery is governed by how the adapter drivers use the BIOS of the VaN. The numbering of each type of adapter is independent of the other. Thus, if you have one encoder and one decoder, each will be card 0 of its type. The order of the adapter slot search is 2, 1, 5, 4, 3.

backup

Specifies that this physical device is to be reserved as a backup device for this VaN. You can specify a maximum of one encoder and one decoder as backup devices.

You may not include a backup device in a logical terminal definition; the only role of a backup device is to wait in reserve in case another similar device fails. If a device does fail, the program that was playing through the failed device is routed automatically to the backup device.

For a decoder, this means that transport stream is sent from the source (usually the ATM adapter) to the backup decoder. For an encoder, it means that the transport stream is taken from the backup encoder and routed to the sink (again, usually the ATM adapter). This backup capability is provided so that installations with high availability requirements can maintain a program across an encoder or decoder outage with minimal interruption.

To take advantage of this feature, you must have appropriate switching equipment and procedures in place to switch analog devices over from the failed adapter to the backup.

virtual

Indicates that this is a virtual device. If you do *not* specify this parameter, the device is assumed to be a physical adapter card, and a *cardNumber* must be specified. See the *IBM 8300 VaN Programming Reference and Guide* for information on virtual device functions.

exename

Is the OS/2 directory and file name of the executable file (usually having a file extension of .exe) that implements the virtual device. If this parameter is specified, AVS will start the executable file via the DosExecPgm system call. The effective command string used to start the program will be:

```
exename devname parmstring
```

Where *exename, devname* and *parmstring* are the corresponding parameter values as specified on the subject device record.

parmstring

The parameter string to be passed to the file identified by the *exename* parameter when it is started. This string may be enclosed in double quotes to allow imbedded blanks, hyphens, and other OS/2 delimiter characters. Note that the when the virtual device executable is started, this parameter string will be prefixed with the device name and a blank.

fileplayer

Indicates that the device is the built-in file player device. This is a virtual encoder that reads a raw MPEG-2 transport stream from a local disk file and makes it available for transmission from a logical terminal.

file recorder

Indicates that the device is the built-in file recorder device. This is a virtual decoder that receives a raw MPEG-2 transport stream and writes it into a local disk file.

mpegfilename

The OS/2 directory path and file name of the MPEG-2 file to b read or written. This name can be overridden later by using the -file parameter on the avsprogram start. command or with the avsprogram record command. If this name is not specified, then the default path and name of D:\$TEMP$.M2T will be used.

videoFormat

The video format parameter may be one of the following choices. The correct one is determined by the capabilities of the audiovisual equipment to which the encoding or decoding device is connected. Consult with the operation manual for your audiovisual equipment to determine the correct setting for this parameter.

MNTSC358 (M)NTSC format using a 3.58 MHz chroma subcarrier. This is the default video format and is the standard format used by video equipment in the United States.

MNTSC443 (M)NTSC format using a 4.43 MHz chroma subcarrier.

NNTSC358 (N)NTSC format using a 3.58 MHz chroma subcarrier.

NNTSC443 (N)NTSC format using a 4.43 MHz chroma subcarrier.

NPAL443 (N)PAL format using a 4.43 MHz chroma subcarrier. This is the predominant video format in most European countries.

NPAL358 (N)PAL format using a 3.58 MHz chroma subcarrier.

MPAL443 (M)PAL format using a 4.43 MHz chroma subcarrier.

MPAL358 (M)PAL format using a 3.58 MHz chroma subcarrier.

Usage

- Each device record must precede all lterm records which contain the device.
- Each device record that defines a virtual device, a fileplayer or a file recorder will cause a separate instance of the device to be started. For file players, file recorders, and for virtual device definitions tha specify an *exeName*, the device instance is started in its own OS/2 process when AVS is started.

Examples: All of the following are valid device definitions:

```
device encoder1 encoder 0
device "Teaching Center: Camera 6"  encoder 1 npal443
device TVMonitor12 decoder 0
device BackupDecoder decoder 2 backup
device "Gary's Cool Virtual Device" decoder virtual c:\gary\vdevice\cool.exe "debug"
device dummyenc encoder virtual c:\ibmavs\samples\avsvenc.exe
device YetAnotherVdev encoder virtual
device playFromDisk fileplayer
device playToDisk file recorder d:\mpegclips\newclip.m2t
```

1. *encoder1* is defined as the first physical Encoder Adapter to be discovered.
2. *"Teaching Center: Camera 6"* is defined as the second physical Encoder Adapter to be discovered. The video format is also set to (N)PAL with a chroma subcarrier of 3.58 MHz.
3. *TVMonitor12* is defined as a the first physical Decoder Adapter to be discovered.
4. *BackupDecoder* is defined as a the third physical Decoder Adapter to be discovered and is reserved as a backup device.
5. *"Gary's Cool Virtual Device"* is defined as a virtual decoder which is implemented in the executable file named c:\gary\vdevice\cool.exe. Since the executable is identified on the device record, AVS will invoke it in its own OS/2 process when AVS is started on the VaN using the following command string:

   ```
   c:\gary\vdevice\cool.exe "Gary's Cool Virtual Device" "debug"
   ```

6. *dummyenc* is defined as a virtual encoder implemented in the executable file named c:\ibmavs\samples\avsvenc.exe. As in the example above, this device will be started at AVS startup with the following command string:

   ```
   c:\ibmavs\samples\avsvenc.exe dummyenc
   ```

7. *YetAnotherVdev* is defined as a virtual encoder. Since no executable name is supplied, the device must be manually started after AVS has started.
8. *playFromDisk* is defined as an IBM file player device. Since no MPEG-2 file path is specified, the default will be used unless it is overridden at some point after AVS has been started.
9. *playToDisk* is defined as an IBM file recorder device. Video streams received by this device will be recorded into the file named

d:\mpegclips\newclip.m2t unless this name is overridden at some point after AVS has been started.

5.1.2.2 The h245 maxsessions Record

This record sets the maximum number of concurrent H.245 control sessions. Since one H.245 session exists on the VaN for each active H.310 terminal connection, this number sets the limit for the number of active connections as well as audiovisual programs that may exist between the VaN and other H.310 terminals at any given time. This number is also used as the maximum number of logical terminals that may be opened at a given time on the VaN.

Format

```
►►──h245maxsessions──nnn──►◄
```

Parameters

nnn
: The nnn code specifies the maximum number of H.245 sessions that can be established by this VaN at one time. This may be a number between 1 to 245, inclusive. If this record is not specified, AVS uses a default value of 20 sessions.

5.1.2.3 The H245 time-out Record

This record specifies the number of seconds each H.245 signaling entity waits for acknowledgements before indicating an error.

Format

```
►►──h245time-out──nn──►◄
```

Parameters

nn The nn code specifies the number of seconds the signaling entity will wait before signaling an error condition. This may be a number between 1 and 30, inclusive. If this record is not specified, AVS uses a default value of 15 seconds.

5.1.2.4 The lterm Record

This record describes a logical audiovisual terminal.

Format

```
►►──lterm──ltermname──atmSelector──devname──┬─────────┬──►◄
                                            └─devname─┘
```

Parameters

ltermname
: specifies the name of this logical terminal. This name is case&hypen.insensitive and may be between 1 to 255 characters long. Since logical terminal names are also IP host names, they must be named according to the syntax rules for host names.

atmSelector
: specifies the selector value of the ATM address (the 20th or least-significant qualifier in the ATM address) for this logical terminal. This must be a number beteen 1 and 127, inclusive. This value identifies the logical terminal on this ATM node and is also used to build the TCP port number on which the logical terminal′s H.245 control sessions will communicate.

devname [*devname*]
: Specifies the names of the devices that are to be part of this logical terminal's configuration. Each of these names must be defined on the device record which precedes this lterm record in the configuration file. Up to 1 input device and 1 output device can be specified.

Usage

- All device names used in each lterm record must be defined in a device record which precedes the lterm record in the configuration file.
- Multiple lterm records may specify the same device names. However, once one of these logical terminals becomes active, the devices it contains are dedicated for use by that terminal. As such the devices are unavailable to any other logical terminal and will remain so until the logical terminal is closed.

Examples: All of the following are valid logical terminal definitions:

```
lterm sendOnlyTerminal 14 encoder1
lterm LoungeMonitor 5 TVMonitor12
lterm StaffConference 63 "Teaching Center: Camera 6" "Teaching Center: Monitor 6"
lterm ClipPlayer 21 playFromDisk
```

1. *sendOnlyTerminal* is a send-only logical terminal with an ATM selector of 14. The encoding device is named *encoder1*.
2. *LoungeMonitor* is a receive-only logical terminal with an ATM selector of 5. The decoding device is named *TVMonitor12*.
3. *StaffConference* is a receive-and-send logical terminal with an ATM selector of 63. The encoding device is named *"Teaching Center: Camera 6"* and the decoding device is named *"Teaching Center: Monitor 6"*.
4. *ClipPlayer* is a send-only logical terminal with an ATM selector of 21. The encoding device is the IBM File Player device named *playFromDisk*.

5.1.2.5 The supervised Record

This record tells AVS that the logical terminals on this VaN may only accept inbound connection requests upon approval by an application program. Therefore, this mode requires that an application program open a logical terminal before it can accept connections.

Format

```
►►──supervised──►◄
```

Usage

- If both this record and the unsupervised record are specified, the last one encountered will be used.
- Applications are notified of incoming connection requests via callback events. Please refer to the *IBM 8300 Video Access Node Programming Reference and Guide* for a complete description of the AVS application programming interface.

5.1.2.6 The tcpportbase Record

This record specifies the base number from which all TCP port numbers are derived. TCP ports are used to establish connections between logical terminals for H.245 sessions. Actual port values are determined by adding the logical terminal′s ATM selector value (from the lterm record) to this base port number.

Format

```
►►──tcpportbase──nnnnn──►◄
```

Parameters

nnnnn
is the base number to be used. This must be a numeric value from 6000 to 65400. If this record is not specified, AVS uses a default value of 10000.

Usage

- All VaNs in the ATM network which will communicate with each other must have the same tcpportbase value.

5.1.2.7 The unsupervised Record

This record tells AVS that the logical terminals on this VaN should accept inbound connection requests (barring any error conditions) when an application program is available to approve or deny the request. Note that when an application program *is* available (in other words, when the logical terminal has been opened by an application program), the application will be consulted to determine whether each connection should be accepted or rejected.

Format

```
►►──unsupervised──►◄
```

Usage

- If both this record and the supervised record are specified, the last one encountered will be used.
- Applications are notified of incoming connection requests via callback events. Please refer to the *IBM 8300 Video Access Node Programming Reference and Guide* for a complete description of the AVS application programming interface.

5.1.2.8 The # Comment Line

The “#” character indicates that any text following the character is to be treated as a comment. Comments may exist on their own lines in the AVS configuration file or may be included on lines with valid configuration records.

Format

```
►►──┬──────────────────────┬──#──comment text──►◄
    └─configuration record─┘
```

Parameters

configuration record
Is any valid AVS configuration record as described in this section.

Is the comment delimiter

comment text
Is the text of the comment

The VaN is shipped with a sample AVS.CFG file, such as the one listed below, for a VaN configured with one encoder and two decoders:

```
#
# Sample configuration file for Video Access Node Audiovisual Services
#
unsupervised
tcpportbase 10000
#
#-----------------------------------------┐
#         Device Definitions              │
#-----------------------------------------│
#    │ Device     │ Device     │ Card     │
#    │  Name      │  Type      │ Number   │
#-----------------------------------------┘
device dec0       decoder      0
device dec1       decoder      1
device enc0       encoder      0
device filePlay   filePlayer   D:\DEM704_5.M2T
device fileRec    fileRecorder
#
#-----------------------------------------┐
#    Local Logical Terminal Definitions   │
#-----------------------------------------│
#    │ Terminal   │ ATM        │ Device   │
#    │  Name      │ Selector   │ Name     │
#-----------------------------------------┘
lterm RAST1         1          enc0,dec0
lterm ROT1          2          dec0
lterm SOT1          3          enc0
lterm ROT2          4          dec1
lterm clipRecorder  5          fileRec
lterm clipPlayer    6          filePlay
```

5.1.3 The AVSTERM.MAP File

The C:\IBMAVS\ETC\AVSTERM.MAP contains the mapping between the terminal names and their ATM addresses. Note that the full 20-byte ATM address is entered here, and the selector byte must be the same one defined for the terminals in the AVS.CFG file.

We suggest that you edit the AVSTERM.MAP file on one of the VaNs and then copy this file to all your other VaNs. This way you can guarantee you will have the same ATM address mapping throughout the network.

A sample AVSTERM.MAP file, such as the one listed below, is shipped with each VaN:

```
#
# Sample terminal address directory file for
# Video Access Node Audiovisual Services
#
# Note: This file is ignored when the '-localtest' option is specified
#
#                                  HN (Hub Number) >┐
#                      ACN (ATM Cluster Number) >┐  │
#              --------------------------------  │  │
#              │ AFI-RDN                       │  │  │          Selector >┐
#              │ ATM subnetwork to which the   │  │  │                    │
#              │ ATM Workgroup switch belongs. │  │  │ -----------------  │
#              │ Must match ATM switch.        │  │  │ │  ATM MAC     │  │
#              │                               │  │  │ │  ADDRESS     │  │
RAST1         39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.83.00.D0.0D.01
ROT1          39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.83.00.D0.0D.02
SOT1          39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.83.00.D0.0D.03
TESTROT       39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.83.00.D0.0D.04
TESTSOT       39.99.99.99.99.99.99.00.00.99.99.01.02.40.00.83.00.D0.0D.05
```

The AVSTERM.MAP edited for the scenarios examined here is listed below:

```
#
# Sample terminal address directory file for
# Video Access Node Audiovisual Services
#
# Note: This file is ignored when the '-localtest' option is specified
#
#                                  HN (Hub Number) >┐
#                      ACN (ATM Cluster Number) >┐  │
#              --------------------------------  │  │
#              │ AFI-RDN                       │  │  │          Selector >┐
#              │ ATM subnetwork to which the   │  │  │                    │
#              │ ATM Workgroup switch belongs. │  │  │ -----------------  │
#              │ Must match ATM switch.        │  │  │ │  ATM MAC     │  │
#              │                               │  │  │ │  ADDRESS     │  │
RAST1         39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.14.01
ROT1          39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.14.02
SOT1          39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.14.03
TESTROT       39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.14.04
TESTSOT       39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.14.05
RAST2         39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.15.01
ROT2          39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.15.02
SOT2          39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.15.03
```

```
RAST3        39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.16.01
ROT3         39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.16.02
SOT3         39.99.99.99.99.99.99.00.00.99.99.01.02.51.12.13.14.15.16.03
MEDIAS1      39.99.99.99.99.99.99.00.00.99.99.01.02.08.00.5A.99.70.BA.00
VDM1         39.99.99.99.99.99.99.00.00.99.99.01.03.40.40.40.40.40.40.00
```

The syntax and conventions used for the AVSTERM.MAP are described in Table 7 on page 117.

Table 7. AVSTERM.MAP Information Contents

Keyword	Description and Parameters
#	Comment line
[no keyword]	Identifies an audiovisual terminal's name and its ATM address. Parameters are: *termName atmAddr # comment* The terminal specified by *termName* is located on the ATM network at the address specified by *atmAddr*. One termName-atmAddr mapping can be included per line. If a terminal resides in a different IP domain, you must specify its fully-qualified IP host name.

5.1.4 Editing the HOSTS File

You must edit the C:\MPTN\ETC\HOSTS file to include all the logical terminals defined in your network. This step must be executed independently if you are using the ATM or the LAN network for the IP control signaling.

We suggest that you edit the HOSTS file on one of the VaNs and then copy this file to all your other VaNs. This way you can guarantee consistent IP address translation throughout the network.

The IP address defined in this file must be the IP address of the interface selected to transport IP traffic: either the ATM or the LAN interface.

You may also include the IP host name defined for your VaN. All the logical terminals defined for the VaN have the same IP address of the VaN interface. The following hosts file is an example for a network of three VaNs, each with a ROT, a SOT and a RAST logical terminal:

```
#
# c:\mptn\etc\hosts
#
9.1.1.3     vanone
9.1.1.3     ROT01
9.1.1.3     SOT01
9.1.1.3     RAST01
9.1.1.4     vantwo
9.1.1.4     ROT02
9.1.1.4     SOT02
9.1.1.4     RAST02
9.1.1.5     vanthree
9.1.1.5     ROT03
9.1.1.5     SOT03
9.1.1.5     RAST03
```

5.1.5 Configuring the IP Parameters for the ATM Interface

You must define IP parameters such as IP address, IP network mask, segment size, maximum transmission unit and ARP server address for the ATM interface on the VaN.

The VaN uses Classical IP over ATM to exchange control messages. For that reason, you need one ATM ARP server in the network, responsible for resolving

ATM addresses to IP addresses. The VaN can to act as an ATM ARP client as well as an ATM ARP server.

The following section detail the configuration steps required in both cases.

5.1.5.1 ARP Configuration

The VaN software includes an extension of the OS/2 Warp V4 TCP/IP support allowing ATM to be used as a transport. The VaN is configured to use this TCP/IP-over-ATM support through the IPATMCFG program. To start this program, enter `IPATMCFG` at an OS/2 command prompt.

You have to choose it the VaN is a client or a server ARP. Only one VaN may be configured as an ARP server. The remaining VaNs must be configured as ARP clients.

You will also have to configure:

- LAN interface: 1
- IP address: 9.1.1.4
- Subnet mask: 255.255.255.0
- ARP server IP address: 9.1.1.3
- ARP server ATM address: 39999999999999000099990102511213141514 00
- Segment size: 4096
- MTU: 4136
- BE peak bit rate: 24

The parameters shown here reflect our scenario. You have to change these values to the ones defined in your environment.

The client configuration requires the IP address and the ATM address of the ARP server.

Reboot the VaN after making any change.

5.1.6 Configuring the IP Parameters for the LAN Interface

The use of the LAN interface is optional, but it may be useful to connect the VaN to an already established LAN for remote applications purposes, such as network management and file sharing.

You may use the built-in Ethernet adapter or an additionally installed token-ring adapter.

To configure the LAN interfaces, execute the following steps:

1 Double-click on the **TCP/IP Configuration (LAN)** icon, at the Main folder (as shown in Figure 51 on page 97). If the TCP/IP Configuration icon is not available, execute it from the command line, by typing `C:\TCPCFG`.

2 The TCP/IP Configuration Panel shows the network parameters of the LAN 0 interface, as illustrated on Figure 60 on page 119.

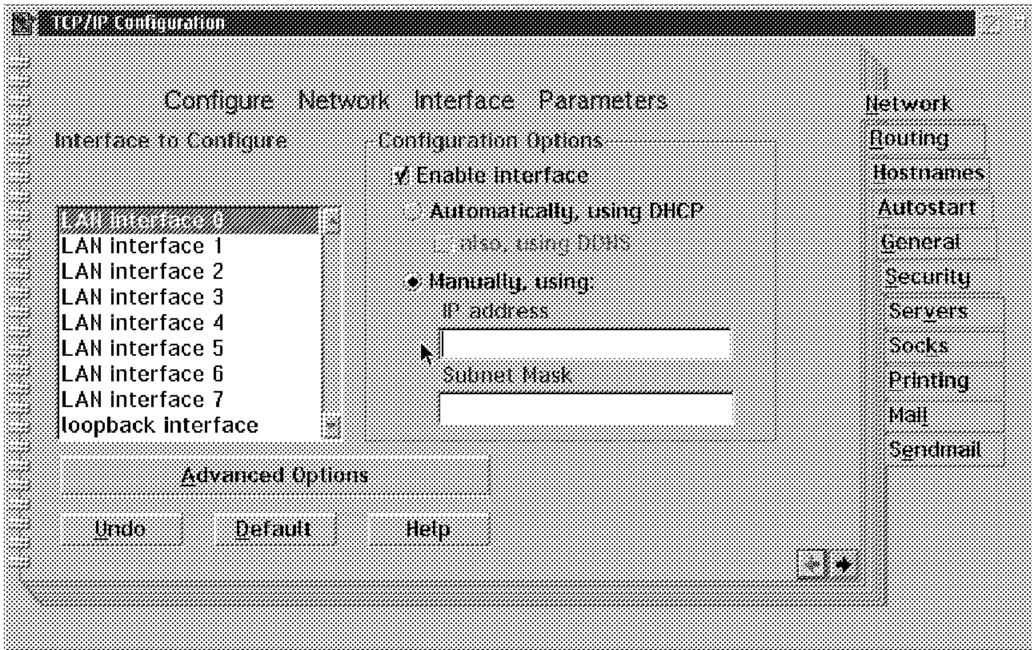

Figure 60. TCP/IP Configuration Panel

3 Click on the **Enable interface** button. Select the **Manually, using:** option and enter the IP address and Subnet Mask information in the corresponding fields.

4 Click on **Hostnames**.

5 On page 1 of the Hostname panel, enter the VaN′s host name. If you wish to use name server services, also enter the local domain name and name server addresses.

6 On page 2 of the Hostname panel, click on **Look through HOSTS list before going to nameserver** option.

7 Close the TCP/IP Configuration Program by double-clicking on the upper left corner of the window. The panel shown in Figure 61 will appear.

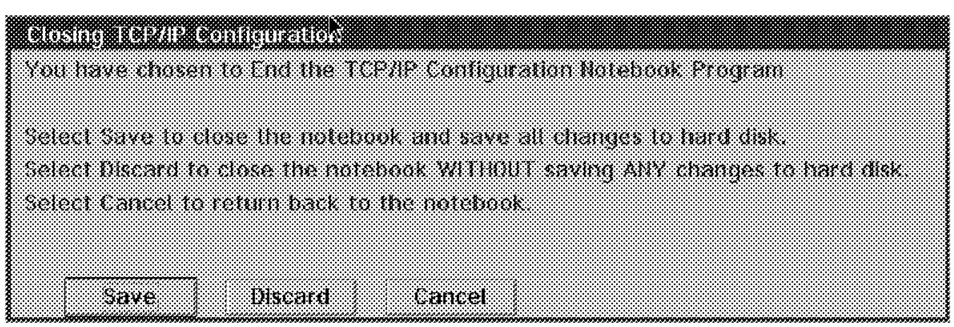

Figure 61. TCP/IP Configuration Message Panel

8 Select **Save** to exit the TCP/IP Configuration Program.

If you are going to connect the VaN to a LAN server, you must edit the file C:\IBMLAN\IBMLAN.INI. The fields, Computername and Domain, come with default values and need to be changed to reflect your specific LAN environment.

5.2 Basic VDM Setup

This section describes the configuration procedure for the Video Distribution Module.

5.2.1 ATM Addressing and the VDM

The VDM module is shipped with the End System Identifier (ESI) portion of the ATM address already burned into the module. This address is taken from a range provided by the IEEE and guaranteed to be unique. You can configure your VDM to use this address or supply your own ESI portion of the ATM address.

The VDM uses a single ESI address to identify the module. If you are using SVCs, use the selector byte portion of the ATM address to identify which of the eight ports on the VDM is being addressed. The correspondence between selector byte values and VDM ports is shown in Table 8.

Table 8. Selector Byte Values per Port

Selector Byte Value	Port
X′00′	1
X′01′	2
X′02′	3
X′03′	4
X′04′	5
X′05′	6
X′06′	7
X′07′	8

When using PVCs, you must specify the virtual path identifier (VPI) and virtual channel identifier (VCI) for the port you are addressing. Correspondence between the VPI/VCI and the ports are shown in Table 9.

Table 9. VPI/VCI Values per Port

VPI/VCI	Port
0/32	1
0/33	2
0/34	3
0/35	4
0/36	5
0/37	6
0/38	7
0/39	8

5.2.2 Connecting to the ATM Backplane

You must attach the VDM to the ATM backplane in order to connect to the ATM network via the hub. Use the following steps to attach the VDM. Refer to the *IBM 8260 Nways Multiprotocol Switching Hub ATM Control Point and Switch Module Installation and User's Guide*, SA33-0326, for details about the commands and problem determination.

At the A-CPSW or 8285 console:

1 Connect to the ATM switch by entering SET MODULE *slot* CONNECTED (where *slot* identifies the lowest numbered slot occupied by the VDM).

2 Enable the ATM port connection by entering SET PORT slot 2 ENABLE UNI (where *slot* identifies the lowest numbered slot occupied by the VDM).

3 Save the ATM switch configuration for the VDM by entering the SAVE MODULE_PORT command.

4 If you want to verify the ATM switch configuration for the VDM backplane connection, enter SHOW MODULE slot VERBOSE (where *slot* identifies the lowest numbered slot occupied by the VDM).

Note

In order to use the VDM with your hub, the hub must be compatible with FPGA version B50 or higher.

5.2.3 VDM Configuration

After you have installed and enabled the VDM on the 8260 or the 8285, you must configure a few parameters. Access to the Operational and the Off Line Service menus of the VDM is done through an ASCII-type terminal connected to its own serial port.

Note

Changes made using the VDM Operational Menu, except for the color bars setting, do not take effect until the VDM is restarted. You can restart the VDM using the **Reset** button on the module, or through the hub console.

The following steps describe the procedures to configure the VDM:

1 Connect your terminal to the VDM RS-232 port and select the **Reset** button on the VDM faceplate. The VDM Off Line Service Menu will appear on your display.

Note

The VDM Off Line Service Menu is only available if there is a terminal connected to the serial port of the VDM when the module is powered-on or reset. Otherwise, the module goes directly to Operational State and the Service Menu is displayed when a terminal is connected at a later time.

```
VDM Off Line Service Menu
 (C) Copyright IBM Corp. 1997. All rights reserved.
-- Type the number of the desired action--
------------------------------------------
    1. Software Upgrade
    2. Change to Operational State
------------------------------------------
OFFLINE==>
```

2 Type 2 to go to the operational state. The VDM Service Menu will appear on your display

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]' is optional argument
SERVICE==>
```

3 Type de to display the ESI address of the VDM module. The default ESI address and the active ESI addresses are displayed.

The default ESI address is the burned-in 6-byte MAC address of the module, taken from a range of addresses assigned by the IEEE. This value cannot be altered.

The active ESI address can be changed. You may want to use a locally administered address (LAA) for your VDM, for easier identification or if your network uses a specific addressing plan.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]'  is optional argument
SERVICE==>  dm
Default ESI: 010203040506   Active ESI: 010203040506
```

4 If you want to use a locally administered address, type ae. The current active ESI will be displayed. Type the new ESI beside the current one and press Enter. You are returned to the Service Menu.

When the VDM is restarted, it will use the new ESI address you have provided.

To revert to the default ESI address, select the **Alter ESI address** option and change the ESI address to all zeros (0).

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]'  is optional argument
SERVICE==>  am

000000000000  112233445566
```

5 If you changed the ESI, ensure you entered the right address by typing de.

6 You must set the operating mode of the VDM: PVC mode or SVC mode. To check what the current VDM mode is, type pvc. The PVC state will be displayed.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]' is optional argument
SERVICE==> pvc
PVC State is: PVC
```

7 If you want to set the VDM to SVC mode, type pvc off.

Note

If you are going to establish a VaN-to-VDM connection, you must use SVC mode.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]' is optional argument
SERVICE==>  pvc off
```

8 Check the current VDM operation mode, by typing pvc.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]' is optional argument
SERVICE==>  pvc
PVC State is: SVC
```

9 If you want to go back to PVC mode, you must type pvc on.

10 Check the connection between the TV and the VDM, using the Color Bars feature. Type cb 1 on to check the connection on port 1. This command generates a color bar image on the TV display.

This feature is specially useful if there is a large number of TVs installed or the coaxial cables used for the TV connection are not easily accessible.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]'  is optional argument
SERVICE==>  cb 1 on
```

11 Type cb. The VDM console will display the color bar state of all ports.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]'  is optional argument
SERVICE==>  cb
Port 1 color bars are on
Port 2 color bars are off
Port 3 color bars are off
Port 4 color bars are off
Port 5 color bars are off
Port 6 color bars are off
Port 7 color bars are off
Port 8 color bars are off
```

12 After you have confirmed proper connection between the VDM ports and the corresponding TVs, turn off the color bar generation. To disable the color bar signal on port 1, type cb 1 off.

13 Check the state of each of the ports on the VDM by selecting **Display Port Status** from the Operational Menu. The possible states are:

Reset State The VDM has just completed a Reset, or the ATM address for this port has been deactivated by the network.

The port could be in this state because its address was refused by the network. The most common cause of this problem is that the address is a duplicate. While the port is in this state, address registration will be requested periodically by the network.

Address Registration State (SVC only)
The VDM is attempting to register the ATM address for this port with the network.

The port could be in this state because a previous address registration request was refused by the network. The most common cause of this problem is that the address is a duplicate.

Ready State (SVC only)
The ATM address for this port has been successfully registered with the network and the VDM is waiting for an SVC call setup.

Connected State
The port is monitoring its assigned VCI for video data.

Disabled State The port is not operational due to a detected hardware error.

After you have configured the VDM, the ATM switches and the source equipment (either a VaN, a MediaSTREAMER or another standards-compliant audiovisual terminal) for all VDM ports, type st to display the current status of the VDM ports.

```
VDM Service Menu
  (C) Copyright IBM Corp. 1997. All rights reserved.
Display Port Status  - st
Display VPD  - vpd
Display ESI Address  - de
Alter ESI Address - ae
Alter PVC mode - pvc [on|off]
Color Bars - cb [ port on|off]

'|'  is choice, 'port' is  1-8, '[]'  is optional argument
SERVICE==> st
Port 1 is in the CONNECTED state
Port 2 is in the CONNECTED state
Port 3 is in the CONNECTED state
Port 4 is in the CONNECTED state
Port 5 is in the CONNECTED state
Port 6 is in the CONNECTED state
Port 7 is in the CONNECTED state
Port 8 is in the CONNECTED state
```

5.3 Configuration Scenarios

Some scenarios involving the VaNs, the VDM and the MediaSTREAMER were implemented to serve as a practical interoperability example. Each of these configurations is explained in details in the following sections.

An ATM network is required as the backbone for these implementations. There are no special requirements for this network and it may consist of various ATM switches. The ATM network used for these scenarios consists of:

One 8260-A17 with:

– ATM Switch/Control Point

- 2-port 155 Mbps ATM Module in slot 2
- Video Distribution Module (VDM) in slot 14

One 8285-00B with:

- Expansion unit
- 2-port 155 Mbps ATM Module in slot 2
- Video Distribution Module (VDM) in slot 3

The ATM addresses of the switches are listed in Table 10.

Table 10. ATM Address of the Switches

ATM Switch	**ATM Address**
8260	39.99.99.99.99.99.99.00.00.99.99.01.03.99.99.99.99.99.99.99
8285	39.99.99.99.99.99.99.00.00.99.99.01.02.99.00.99.00.99.00.01

The 8260 and the 8285 are connected through one 155 Mbps SSI link, using multimode fiber.

The ATM ports used for the VaN, VDM and the MediaSTREAMER were configured as UNI ports, using either 155 Mbps multimode fiber or 25 Mbps UTP.

A complete description of the module and port configuration on the 8260 is listed below as a response to the Show commands issued at a terminal connected to the 8260. The ports on the 8285 are configured with the same characteristics.

```
8260ATM> show module 2 verbose

Slot Install Connect Operation General Information
-------------------------------------------------------------------------------
  2     Y       Y        Y     8260 ATM 2 Ports LAN 155 Mbps Module


status: connected / hardware okay
        enable  / Normal
P/N:51H3635 EC level:E28056 Manufacture: VIME
Operational FPGA version : B50
     Backup FPGA version : 7

     Type  Mode     Status
------------------------------------------------------------------------------
 2.01: SSI enabled  UP-OKAY
 2.02: UNI enabled  UP-OKAY
```

```
8260ATM> show port 2.1 verbose


     Type  Mode     Status
-------------------------------------------------------------------------------
 2.01: SSI enabled  UP-OKAY

SSI Bandwidth        : 155000 kbps
VPI.VCI range        : 15.1023 (4.10 bits)
Connector            : SC DUPLEX
Media                : multimode fiber
Port speed           : 155000 kbps
Remote device is active
IX status            : IX OK

Scrambling mode      : frame and cell
Clock mode           : internal
```

```
8260ATM> show port 2.2 verbose


     Type  Mode     Status
-------------------------------------------------------------------------------
 2.02: UNI enabled  UP-OKAY

Signaling Version   : with ILMI
Flow Control         : On
VPI.VCI range        : 15.1023 (4.10 bits)
Connector            : SC DUPLEX
Media                : multimode fiber
Port speed           : 155000 kbps
Remote device is active
IX status            : IX OK

Frame format         : SONET STS-3c
Scrambling mode      : frame and cell
Clock mode           : internal
```

5.3.1 VaN-to-VaN Configuration

The purpose of this configuration is to set up an audiovisual transmission between two video access nodes.

The VaNs are called VaN1 and VaN2. There will be an audiovisual transmission from VaN1 to VaN2, as well as another transmission in the opposite direction. Figure 62 on page 129 illustrates the network configured for this example.

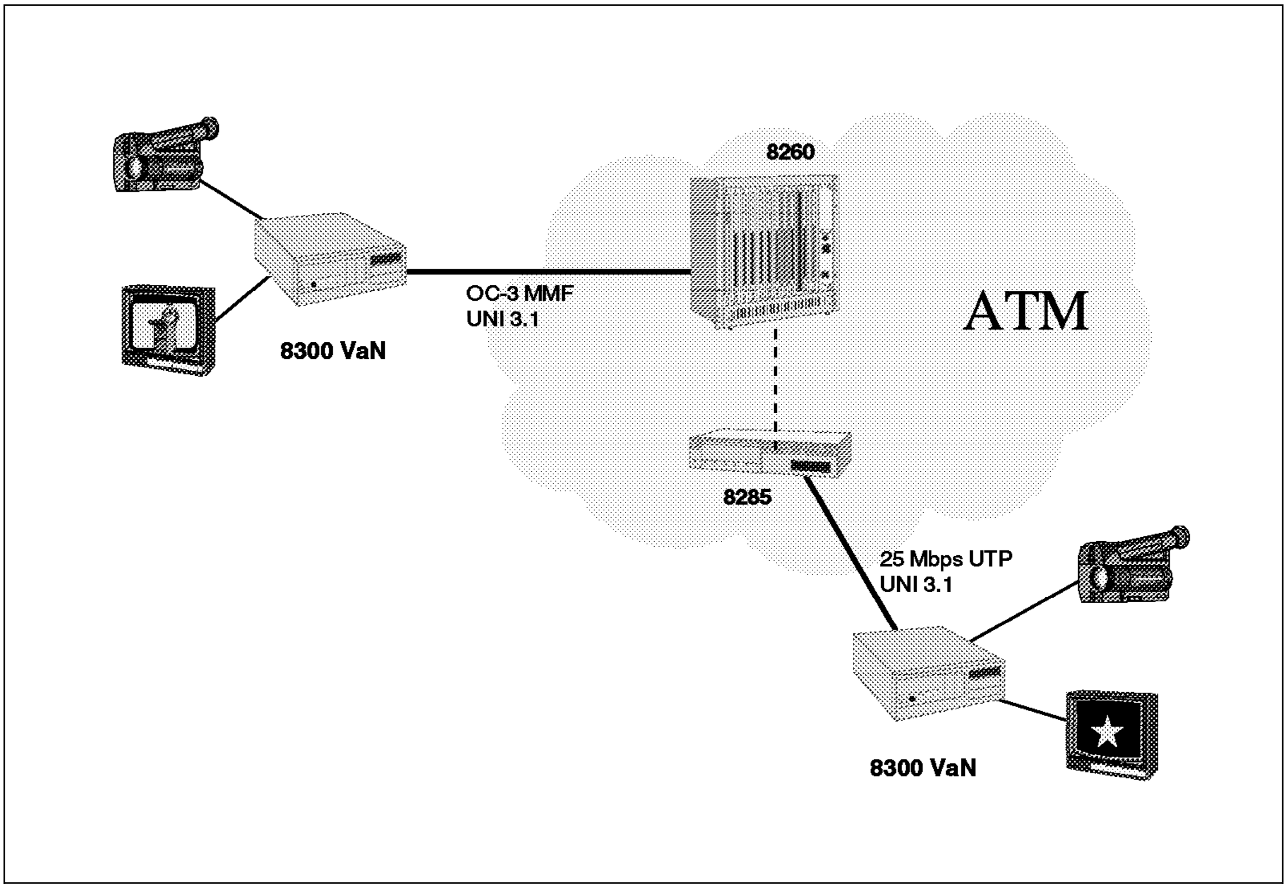

Figure 62. VaN-to-VaN Configuration

The assumption is that the the ATM network has been configured and the ports are enabled.

VaN1 is connected to the ATM network (port 2.2 on the 8260) through a UNI 3.1 link, via 155 Mbps multimode fiber.

VaN2 uses a 25 Mbps UTP connection, with ATM UNI 3.1 signaling, to reach the ATM network (port 1.1 on the 8285).

Both video access nodes have one encoder card and one decoder card. IP support is not configured for their Ethernet adapters.

Table 11 shows a more detailed listing of the VaN settings.

Table 11. Configured VaN Parameters

Parameter Description	**Source**	**Destination**
Name	VaN1	VaN2
Selector Byte - Logical Terminal	01 - RAST01 02 - SOT01 03 - ROT01	01 - RAST02 02 - SOT02 03 - ROT02
ATM Address	39.99.99.99.99.99.99.00.00.99.99. 01.02.51.12.13.14.15.14	39.99.99.99.99.99.99.00.00.99.99. 01.02.51.12.13.14.15.15
IP Address - Network Mask	9.1.1.3 - 255.255.255.0	9.1.1.4 - 255.255.255.0
ATM ARP Role	Server	Client

After the ATM network is setup and the VaNs are configured according to the parameters above and the procedures described in 5.1, "Basic VaN Setup" on page 95, execute the following steps to initialize AVS and start the audiovisual programs:

1 Connect a VCR camera to the encoder card and a TV set to the decoder card of each VaN. Turn on the TVs and the cameras.

2 Start **Audiovisual Services** by double-clicking on its icon, located in the IBM VaN Audiovisual Services folder (see Figure 63).

This must be executed at VaN1 *and* VaN2.

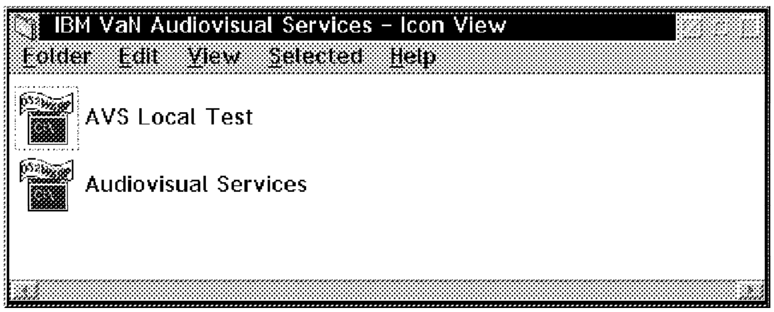

Figure 63. IBM VaN Audiovisual Services Panel

3 The display shows the various status messages and stops with the final message:

`Audiovisual Services are initialized and ready.`

This means that the AVS is active and the VaN can start the audiovisual programs. Do not close this window. The AVS must be active at all times.

At this time the decoders are reset and the TVs will display the color bar signal.

4 To start transmission from VaN1, open another OS/2 window at VaN1 and type `avsuser SOT1 ROT2`.

You will see the message:

`Program started successfully. Press any key to stop the program...`

You can check the messages issued by the VaN at the AVS window. At VaN1, where the avsuser comand was executed, you will read:

`Informational: Encoder card 0 in On-Card Memory mode.`

The AVS window at the destination node (VaN2) will display the message:

`Informational: MPR Decoder 0 started using audio PID 0x1063, video PID 0x1062.`

The video and audio of the camera connected to VaN1 will be displayed at the TV connected to VaN2.

5 To start transmission from VaN2, open another OS/2 window at VaN2 and type `avsuser SOT2 ROT1`.

You will see messages similar to the ones displayed before.

The video and audio of the camera connected to VaN2 will be displayed at the TV connected to VaN1.

5.3.2 VaN-to-VDM Configuration

The topology of the network setup for this configuration is shown on Figure 64.

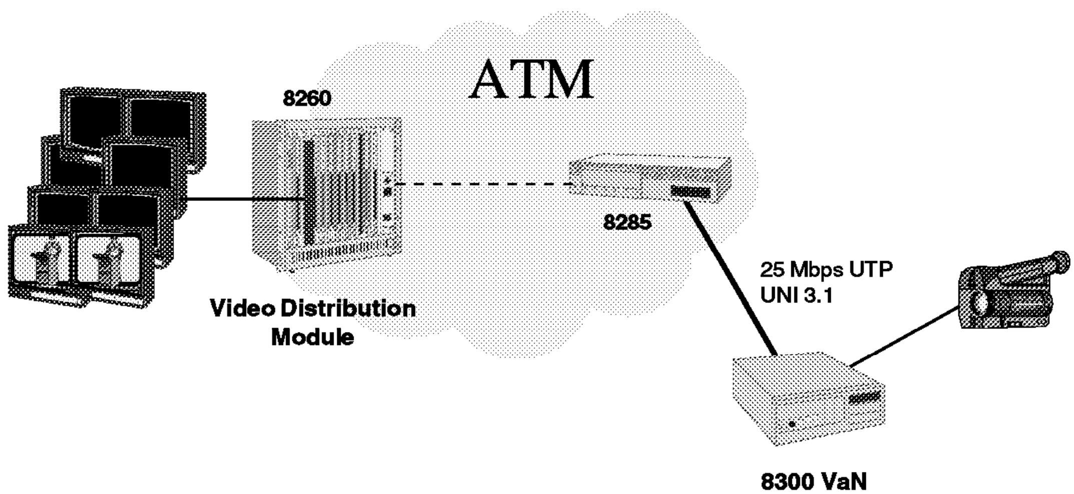

Figure 64. VaN-to-VDM Configuration

At the time this book was written, this interoperability scenario was not possible. Details on the configuration of this topology will be included in the next update of this book.

5.3.3 MediaSTREAMER-to-VDM Configuration

The purpose of this configuration is to set up an audiovisual transmission between the MediaSTREAMER and the VDM installed in the 8285 expansion unit.

An MPEG-2 encoded video file stored on the MediaSTREAMER is played on the MediaSTREAMER and is displayed on eight TV sets connected to the VDM. The file, demo704, played in this example for all eight ports is the same, but different files can be played for each port.

Figure 65 on page 132 illustrates the topology of the network configured for this example.

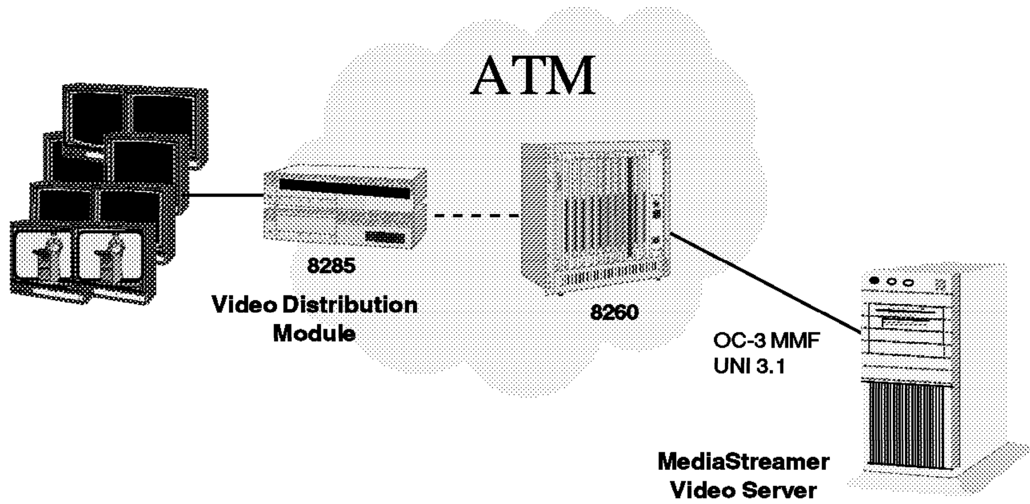

Figure 65. MediaSTREAMER-to-VDM Configuration

The VDM and the MediaSTREAMER were configured to operate in PVC mode. For this reason, eight PVCs were created from the ATM port connected to the MediaSTREAMER to the VDM port.

The PVCs were defined through the 8260 console port, following the PVC-to-VDM port mapping defined in Table 9 on page 120.

It is important to notice that the ATM network may consist of any ATM switch and in any quantity and the TVs could be connected on other VDM ports available on the network, as long as the corresponding PVCs are created on the ATM switches.

The basic assumptions for the ATM network are:

- The PVC endpoints must be:
  - 155 Mbps ATM port on the MediaSTREAMER, configured as UNI 3.1 (port 2.2 on the 8260).
  - VDM port (port 3.2 on the 8285).

  Note

  The VDM occupies two slots on the 8260 or 8285. When configuring the VDM you must refer to the lower slot number being used by the module. You must also refer to all the eight audio and video outputs on the VDM as port number 2. The audio and video outputs are selected through specific VCI values.

- You must configure a PVC service, using AAL5.
- The network must guarantee the bandwidth defined for the created PVCs, even during high traffic levels.

The Set pvc command used to create the PVCs on the 8260 in the tested configuration is shown below:

```
8260ATM> set pvc
Enter local port: 2.2
Enter PVC id: 9
Enter remote port: 3.2
Enter remote hub number: 2
Enter call type: channel_point_to_point
Enter local VPI: 0.
Enter local VCI: 32
Enter remote VPI: 0.
Enter remote VCI: 32
Enter quality of service: reserved_bandwidth
Enter bandwidth in kbps: 15000

PVC set and started.
```

After creating the PVCs, check their status, by typing show pvc all.

```
8260ATM> show pvc all

       Local end point       ! Remote end point        !
-----------------------------+-------------------------+
 Port   id   type   Vpi/Vci  ! Port Vpi/Vci   HNb Party!   role  !QOS! Status
-----------------------------+-------------------------+---------+---+---------
 2.02     1 PTP-PVC   0/32   ! 3.02   0/32      2     0! Primary ! RB!Active
 2.02     2 PTP-PVC   0/33   ! 3.02   0/33      2     0! Primary ! RB!Active
 2.02     3 PTP-PVC   0/34   ! 3.02   0/34      2     0! Primary ! RB!Active
 2.02     4 PTP-PVC   0/35   ! 3.02   0/35      2     0! Primary ! RB!Active
 2.02     5 PTP-PVC   0/36   ! 3.02   0/36      2     0! Primary ! RB!Active
 2.02     6 PTP-PVC   0/37   ! 3.02   0/37      2     0! Primary ! RB!Active
 2.02     7 PTP-PVC   0/38   ! 3.02   0/38      2     0! Primary ! RB!Active
 2.02     8 PTP-PVC   0/39   ! 3.02   0/39      2     0! Primary ! RB!Active
 3.02  1001 PTP-PVC   0/32   ! 2.02   0/32      3     0!Secondary! RB!Active
 3.02  1002 PTP-PVC   0/33   ! 2.02   0/33      3     0!Secondary! RB!Active
 3.02  1003 PTP-PVC   0/34   ! 2.02   0/34      3     0!Secondary! RB!Active
 3.02  1004 PTP-PVC   0/35   ! 2.02   0/35      3     0!Secondary! RB!Active
 3.02  1005 PTP-PVC   0/36   ! 2.02   0/36      3     0!Secondary! RB!Active
 3.02  1006 PTP-PVC   0/37   ! 2.02   0/37      3     0!Secondary! RB!Active
 3.02  1007 PTP-PVC   0/38   ! 2.02   0/38      3     0!Secondary! RB!Active
 3.02  1008 PTP-PVC   0/39   ! 2.02   0/39      3     0!Secondary! RB!Active
```

The only configuration required by the VDM is that it must be in PVC mode. Refer to 5.2, "Basic VDM Setup" on page 120 for the detailed configuration procedure.

The MediaSTREAMER requires a script file to issue the commands to set up and start the video file transmission. The script file VELATEST18.DAT was used to initiate the playback of a video file from the MediaSTREAMER connected to an ATM port on the 8260 to the eight VDM ports installed in the 8285.

Here is the listing of the script file used:

```
#
# Title:  velatest18.dat
# Author: virgil moffett
# Test: vela + 8 vdm port test to play 8 streams of the video file demo704 over ATM
#
Verbose\1\
Init\\\\\\
#Interactive\1\
OpenSession\\\ag0\\
RegisterCallBack\1\MS_ALL_EVENTS\\
OpenPort\1\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:32\\\\\\\\\\\\\\\\\
OpenPort\2\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:33\\\\\\\\\\\\\\\\\
OpenPort\3\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:34\\\\\\\\\\\\\\\\\
OpenPort\4\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:35\\\\\\\\\\\\\\\\\
OpenPort\5\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:36\\\\\\\\\\\\\\\\\
OpenPort\6\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:37\\\\\\\\\\\\\\\\\
OpenPort\7\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:38\\\\\\\\\\\\\\\\\
OpenPort\8\MS_NET_CONNECT_FOR_OUTPUT\5500000\atm\atm0\MS_NET_ATM\MS_PROTO_AAL5\\
AF_NDD\CONN_PVC\\\\\\\\\\\\\\\AF_NDD\CONN_PVC\\\\\\\0
:39\\\\\\\\\\\\\\\\\
OpenPlayStream\1\1\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\2\2\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\3\3\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\4\4\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\5\5\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\6\6\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\7\7\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
OpenPlayStream\8\8\demo704\1\1\MS_PAUSE\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\1\111\\\
Play\1\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\2\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\3\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\4\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\5\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\6\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\7\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Play\8\1\MS_NPT_POS\0\0\MS_NPT_POS\0\0\1\1\\
Interactive\1\
#CloseStream\1\\
ClosePort\1\\
ClosePort\2\\
ClosePort\3\\
ClosePort\4\\
ClosePort\5\\
ClosePort\6\\
ClosePort\7\\
ClosePort\8\\
CloseSession\\
```

The MediaSTREAMER commands executed by this script file perform the following tasks:

- Opens an audiovisual session
- Configures eight logical ports with their operating characteristics, such as:
  - Output ports.

  - Transmission bit rate of 5.5 Mbps.
  - Use of ATM interface (atm0 adapter is a 155 Mbps MMF adapter).
  - Use of AAL5 as the ATM Adaptation Layer.
  - PVC mode.
  - The specific VPI and VCI used by each port. All ports use VPI=0 and the VCI ranges from 32 (port 1) to 39 (port 8).
- Plays an MPEG-2 file, called demo704.
- Closes the logical ports and the audiovisual session.

The following command executes the script file at the MediaSTREAMER:

```
apiDriver VELATEST18.DAT
```

5.3.4 VaN-to-MediaSTREAMER Configuration

The topology of the network setup for this configuration is shown on Figure 66.

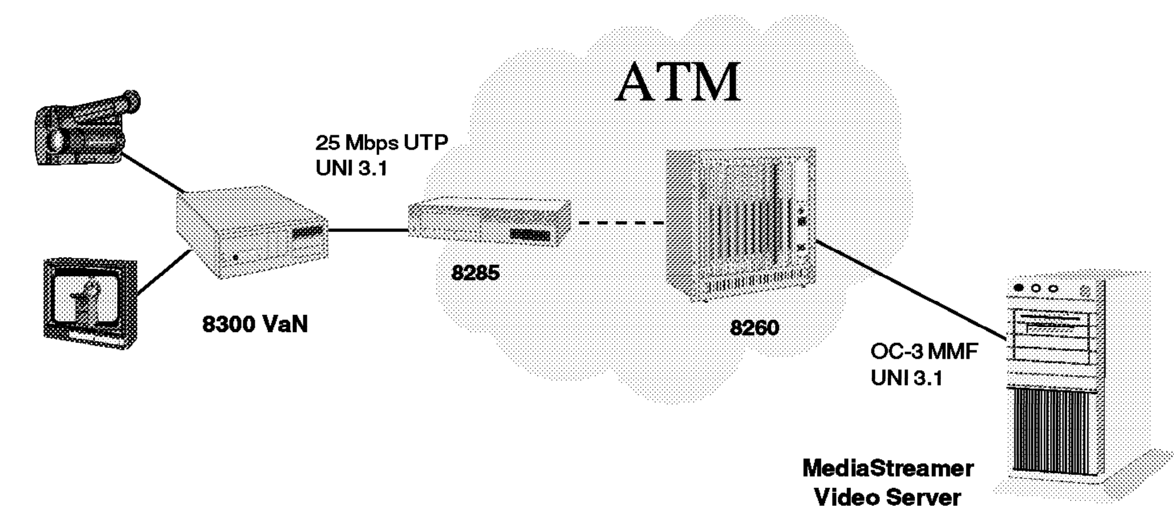

Figure 66. VaN-to-MediaSTREAMER Configuration

At the time this book was written, this interoperability scenario was not possible. Detailed configuration will be included in the next revision of this book.

Chapter 6. Audiovisual Services

The IBM Video Access Node (VaN) is controlled by a software component called Audiovisual Services (AVS). This software runs on IBM Operating System/2 (OS/2) V4. All of the required operating software, which includes AVS as well as OS/2, is preinstalled on the VaN.

The AVS software is based on the International Telecommunications Union (ITU) H.310 standard for broadband audiovisual communication and terminals. The AVS software:

- Presents the VaN to the ATM network according to protocols specified in the H.310 standard. A VaN can be presented as one or more H.310 terminals, depending on local configuration parameters.
- Manages the audiovisual resources of the VaN.
- Manages control flows between the VaN and other audiovisual terminals.
- Controls and monitors the transmission of audiovisual data between the VaN and other audiovisual terminals.

AVS provides three types of user interfaces:

- An application programming interface (API), described in this manual.
- A command-line interface that allows an operator to monitor and control the operation of a VaN. This interface is described in the *Video Access Node Setup and User's Guide*.
- Network management using the Simple Network Management Protocol (SNMP). This interface also is described in the *Video Access Node Setup and User's Guide*.

This chapter provides definitions of concepts and a description of types of communication between terminals.

6.1 VaN Terms and Concepts

A VaN transmits, receives, or both transmits and receives encoded audiovisual data. Transmission of this data occurs between the VaN and another audiovisual terminal, which resides in the same VaN or is physically separate. If the other terminal resides in the same VaN, the encoded audiovisual data never leaves that VaN. If the other terminal does not reside in the same VaN, the audiovisual data is carried over the ATM network. Various types of configuration options allow the VaN to be configured for different environments. Basic concepts and terms are presented in this section to describe these configuration options. Use Figure 67 on page 138 as a reference for this section.

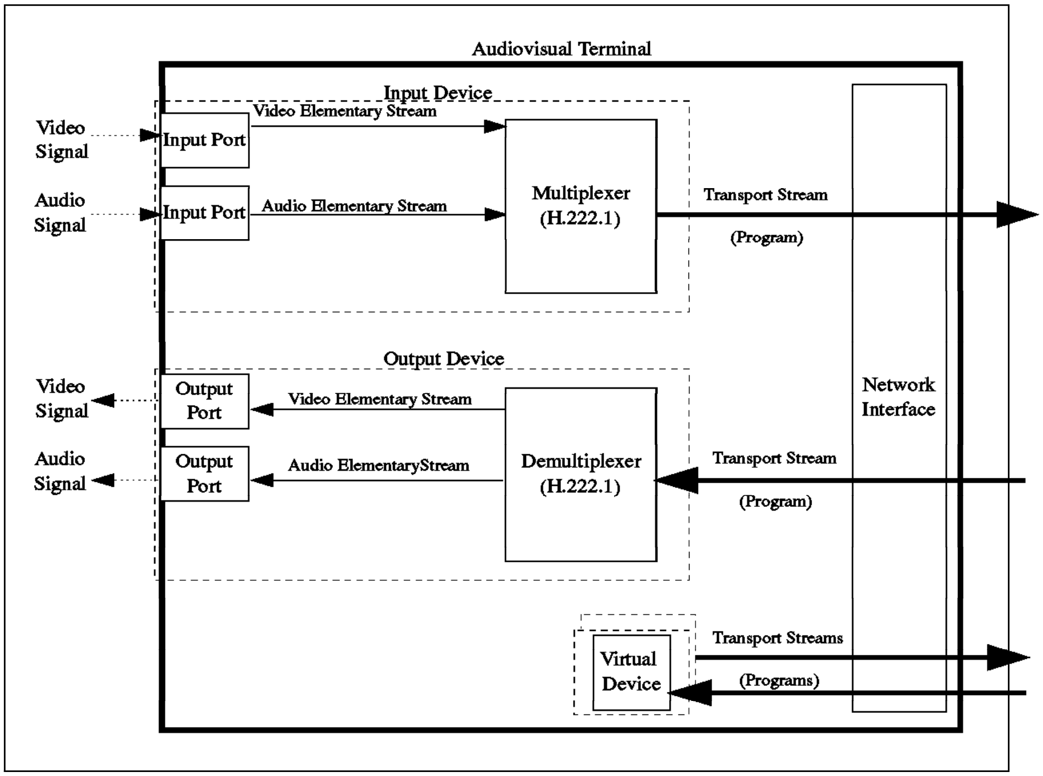

Figure 67. VaN Concepts and Terms

- **Audiovisual Terminal**: An audiovisual terminal, or simply a terminal, contains an ATM network connection, one or more audiovisual devices, and software to manage the resources of the terminal.

  The H.310 standard defines three categories of terminals:

  - Receive-only terminals (ROTs) only receive encoded audiovisual data as input
  - Send-only terminals (SOTs) only send encoded audiovisual data as output
  - Receive-and-send terminals (RASTs) are capable of both operations

- **Audiovisual Devices**: Each VaN can support multiple audiovisual input and output devices. When used in reference to the VaN, an audiovisual device produces or consumes a single transport stream encoded in MPEG-2. This transport stream carries a single set of time-related components, which constitute a program. Because each device can support only a single program, all devices are considered to be unidirectional (input or output).

  A device usually corresponds to either an encoder adapter (produces a transport stream) or a decoder adapter (consumes a transport stream), but the VaN's logical view of devices makes no assumptions about physical configuration or the makeup of the system's devices. It is possible for a device to be implemented completely in software. Such a device is called a virtual device.

- **Ports**: Devices contain one or more ports. When used in reference to the VaN, a port is a part of an audiovisual device that produces or consumes one or more MPEG-2 elementary streams. These elementary streams are the components which, when multiplexed together according to MPEG-2 formats, constitute a program.

A simple example of such a port on an encoder adapter would be a simple video input jack. The device element that links a program and its component elementary streams is the multiplexer (for an encoder adapter) or demultiplexer (for a decoder adapter). In practice, the number of physical connectors (such as video or audio jacks) is not limited to one per port. Rather, ports can consist of multiple physical connectors that often support different formats of the same type of signal (for instance, composite video and S-video). The key concept is that each of these physical connectors is routed to the same elementary stream encoding or decoding component. It is up to the adapter and its device drivers to determine how the multiple physical connectors are supported. However, the use of the local device control APIs allows an application to manipulate these controls with some degree of device independence.

Given these definitions, the job of an input device (an encoder adapter) is to combine the elementary streams from the various input ports into a single program. An output device (a decoder adapter) accepts an input program and demultiplexes it into its component elementary streams, which are then passed along to the output ports on that device. Everything that happens to the program from the time it is generated by an input device to the time it is processed by an output device is the responsibility of the audiovisual terminals that contain the devices.

According to the definitions presented thus far, the VaN itself is a terminal by its own attributes. However, the VaN is designed to contain multiple instances of similar devices (more than one encoder and more than one decoder). H.310 implies that a terminal will have at most one input and one output device, although it does not prohibit terminals that exceed this number. The need exists to be able to group devices so that they fit the H.310 model, that is, having no more than one encoder and one decoder. Logical terminals are used to fulfill this need.

- **Logical Terminal**: A logical terminal is a logical grouping of one or more devices on a VaN that appear to the ATM network as independent H.310 terminals.

  Logical terminals enable the resources of a VaN to be presented to the network in a wide variety of configurations. Because the same physical device can be configured into multiple logical terminals, any combination of the possible configurations can be presented to the network simultaneously. When a device becomes active through a logical terminal, any new request to access that device through another logical terminal will fail with an indication that the resource is not currently available.

  Figure 68 on page 140 illustrates the configuration of a VaN into multiple logical terminals. In this example, the VaN has three devices and is configured into three logical terminals. Note that LTerm A and LTerm B cannot both be active at the same time.

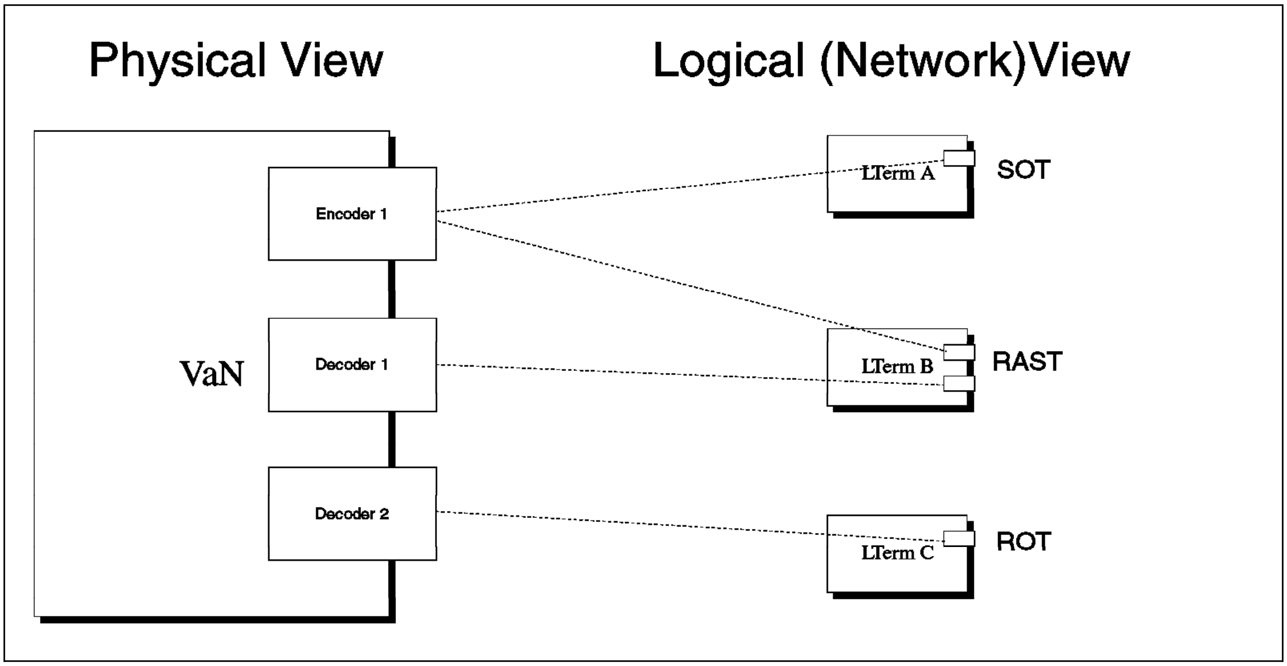

Figure 68. Mapping of Logical Terminals onto a VaN

Logical terminals are defined in the VaN configuration file and are created by AVS software running on the VaN at AVS startup. The logical terminals are accessible from the network from that time until AVS is shut down.

6.2 Communication between Terminals

There are three basic types of communication between a VaN and other audiovisual terminals:

- **Audiovisual channels** carry the actual audiovisual program data between terminals. Each channel is implemented with a dedicated ATM virtual circuit and carries a single program over an MPEG-2 transport stream (as specified by H.222.1) in each direction. Audiovisual channels exist regardless of whether the partner terminal is compliant with H.310 or ATM Forum VoD.
- **Control sessions** send H.245 control information between peer H.310 terminals and are thus used only when communicating with other H.310 terminals. Because the VaN presents itself to the network using logical terminals, control sessions are established on a logical terminal basis. A single control session exists between a given logical terminal and each remote H.310 terminal with which it is communicating, regardless of the number of audiovisual channels that might exist between them. In the VaN, control sessions exist over a TCP/IP connection.
- **Application flows** are communications carried out between an application program running on a VaN and another application program running on a remote terminal. The VaN system software has no direct involvement in such communication. It is up to you to choose an appropriate network transport for these flows. Because TCP/IP is a required component of the VaN software platform, it is anticipated that most application programmers will choose TCP/IP or UDP/IP as their transport.

Figure 69 on page 141 illustrates the various types of communication between a VaN and remote terminals.

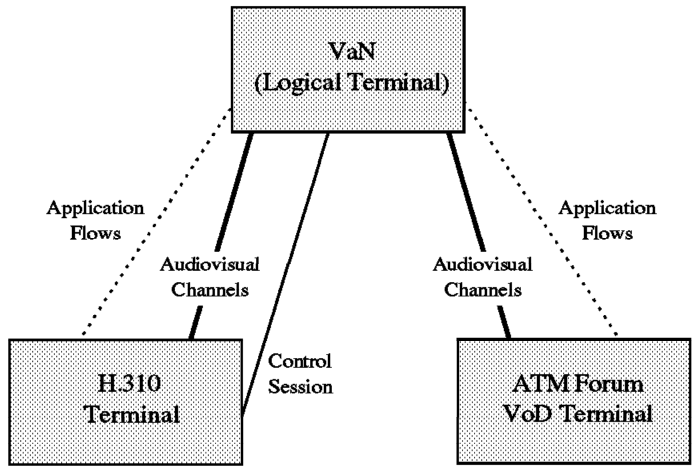

Figure 69. Types of Communication between a VaN and Other Audiovisual Terminals

6.3 Audiovisual Services Commands

This chapter describes commands supported by the VaN Audiovisual Services (AVS) software. As used in this chapter, an *audiovisual terminal*, or simply a *terminal*, contains an ATM network connection, one or more audiovisual devices, and software to manage the resources of the terminal. An *audiovisual device* is an input or output device that either produces or consumes a single transport stream encoded in MPEG-2. In the context of the VaN, a *port* is a part of an audiovisual device that produces or consumes one or more MPEG-2 elementary streams. It is these elementary streams, when multiplexed together according to MPEG-2 formats, that constitute a transport stream, or *program*.

A *logical terminal* is a logical grouping of one or more devices on a VaN that appear to the ATM network as independent H.310 terminals.

6.4 Commands

You can use commands to get information regarding configuration, diagnostics, and operational statistics. Using the commands described in this chapter, you can:

- Start and stop the Audiovisual Services (AVS) software
- Perform basic operational control of installed hardware
- Dynamically update audiovisual network configuration parameters
- Control local VaN resources
- Compile configuration information, operational statistics, and diagnostic information

Commands for performing these tasks are summarized in Table 12 on page 142 and are described in more detail following the table.

Table 12. AVS Commands Summary

Command	Function
avsprogram	Controls transmission of audiovisual programs
avsquery	Displays resource and configuration information
avstermmap	Dynamically updates audiovisual terminal ATM address mapping
avsstart	Starts AVS software
avsstop	Stops AVS software
avstrace	Activates, deactivates, or modifies AVS tracing

6.4.1 The avsprogram Command

This command allows you to modify, start, or stop audiovisual transmissions.

Format

```
►►──avsprogram───┬─addparty──parameters───┬──►◄
                 ├─dropparty──parameters──┤
                 ├─pause──parameters──────┤
                 ├─record──parameters─────┤
                 ├─resume──parameters─────┤
                 ├─start──parameters──────┤
                 └─stop──parameters───────┘
```

Parameters

addparty
: Adds one or more new receivers to a currently-transmitting unidirectional program.

dropparty
: Drops one or more receivers from a currently transmitting unidirectional program.

pause
: Temporarily suspends playback of a locally-stored audiovisual program.

record
: Prepares a logical terminal configured with a **filerecorder** decoder device to record an inbound audiovisual program.

resume
: Resumes the playback of a locally-stored audiovisual program after a pause.

start
: Initiates the transmission of an audiovisual program from a local logical terminal to another audiovisual terminal (either local or remote).

stop
: Terminates the transmission of an audiovisual program.

Each form of the avsprogram is described below.

6.4.2 The avsprogram addparty Command

This command adds one or more new receivers to a currently transmitting itting unidirectional audiovisual program. Issue this command at the transmitting terminal.

Format

```
                                        ┌────────────────────────────────┐
                                        │  ┌─-auto─┐   ┌──────────┐     │
►►──avsprogram addparty──sourceLterm────▼──┤       ├───▼─destTerm─┴─────┴───►◄
                                           ├─-h310─┤
                                           └─-vod──┘
```

Parameters

sourceLterm
The name of the local logical terminal from which the program is being transmitted.

Destination List
This is a list of destination terminal names and optionally, information about how to connect to them. One or more such lists may be specified. Each list is grouped by the connection type, which is indicated by the connection type qualifier that precedes the list of terminal names. The first destination list may be specified with or without a preceding connection type qualifier. If no qualifier is specified on the first list, -auto is assumed. All remaining lists must begin with a connection type qualifier. Multiple lists with similar connection type parameters are allowed.

-auto Indicates that the connection type is to be determined automatically. For each terminal in the sublist, AVS will first attempt to connect to the terminal using H.310 protocols. If the connection fails, AVS will then attempt to connect in ATM Forum Video-on-Demand (VoD) mode. If this connection type is used, the user does not have to know the exact type of each destination terminal. However, connecting in this mode causes a significant increase in the overall connection time for the addparty command.

If no connection qualifier is specified on the first destination sublist, -auto is assumed.

-h310 Tells AVS to connect to the terminals listed between this parameter and the end of the command line or a later -vod parameter using H.310 protocols. Use this option to reduce the processing and connection time for the addparty command when adding other VaNs to the transmission. Note, however, that use of this parameter requires you to know the connection type of the destination terminals.

-vod Tells AVS to connect to the terminals listed between this parameter and the end of the command line or a later -h310 parameter using ATM Forum Video-on-Demand (VoD) protocols. Use this option to reduce the processing and connection time for the addparty command when adding IBM MediaSTREAMERs, IBM Video Distribution Modules, or other ATM Forum VoD terminals to a transmission. Note, however, that use of this parameter requires you to know the connection type of the destination terminals.

destTerm Names of one or more audiovisual terminals to be added as receivers of the current program.

All of the following are valid destination lists:

```
ArlingtonHS QuincyHS AltonHS
ArlingtonHS -vod QuincyHS -h310 AltonHS
- vod ArlingtonHS QuincyHS AltonHS
- vod ArlingtonHS -h310 QuincyHS -vod AltonHS
- h310 ArlingtonHS -auto QuincyHS AltonHS
```

1. The first example defaults to the -auto connection for all listed terminals.
2. The second example defaults to the -auto connection for ArlingtonHS, uses a VoD connection to QuincyHS, and an H.310 connection to AltonHS.
3. The third example shows that the same connection qualifier may be used in multiple sublists. Here, both ArlingtonHS and QuincyHS use VoD connections.
4. The fourth example shows the use of the -auto qualifier.

Usage

- This command may only be used after an audiovisual program transmission has been initiated by using the **avsprogram start** command.
- More than one avsprogram addparty command may be issued against the same audiovisual program.
- You may not add new receivers to a bidirectional audiovisual program (a program that was initiated with the -b option of the avsprogram start command).
- At most, one receiver of an outbound audiovisual program may be another logical terminal on the local VaN. Therefore, if there is currently no local receiver of the transmission, then one may be added by using the **avsprogram addparty**. If a local receiver already exists, then all of the terminals listed in the *destTerm* list must be remote audiovisual terminals connected through the ATM network.
- The addparty command will be attempted for every listed destination terminal. If errors occur for one or more of those terminals, an appropriate error message will appear and processing will continue on to the next terminal in the list.

Examples: This command successfully adds three receivers to the program currently being transmitted from the local logical terminal named *scienceLab*. Since a connection type is not listed, each will be tried using an H.310 connection and then, if necessary, a VoD connection:

```
[C:\]avsprogram addparty scienceLab classroom3 classroom5 -vod classroom7
[C:\]
```

6.4.3 The avsprogram dropparty Command

This command drops one or more receivers from a currently-transmitting unidirectional audiovisual program. Issue this command at the transmitting terminal.

Format

```
                                        ┌─────────┐
                                        ▼         │
►►──avsprogram dropparty──sourceLterm────destTerm─┴──►◄
```

Parameters

sourceLterm
The name of the local logical terminal from which the program is being transmitted.

destTerm
The list of terminal names to be dropped from the transmission. These terminals must have been previously specified as receivers by usin using either the **avsprogram start** or the avsprogram addparty command.

Usage

- More than one avsprogram dropparty command may be issued against the same audiovisual program.
- If the last receiver of the transmission is dropped, then transmission from the source is terminated also.
- You may not drop the receiver from a bidirectional audiovisual program (a program that was initiated with the **-b** option of the avsprogram start command).

6.4.4 The avsprogram pause Command

This command temporarily suspends playback of a locally-stored audiovisual program through a fileplayer device. When a program is paused, all ATM connections remain active, but the transmission of audiovisual data to the destination terminal is stopped. Playback may be resumed by issuing **avsprogram resume**.

Format

```
►►──avsprogram pause──sourceLterm──►◄
```

Parameters

sourceLterm
: The name of the local logical terminal from which the recorded program is being transmitted. The encoding device in this terminal must be of the type, fileplayer.

Usage:

The program being paused must have been started by using the avsprogram start command.

6.4.5 The avsprogram record Command

This command allows you to specify the name of the file into which an inbound audiovisual program is to be recorded and ensures that the logical terminal is available for recording. This command does not initiate an audiovisual transmission.

Format

```
                                   ┌─D:\$TEMP$.M2T─┐
►►──avsprogram record──localLterm──┼───────────────┼──►◄
                                   └─mpegfilepath──┘
```

Parameters

localLterm
The name of a local logical terminal which contains a decoding device of the type, filerecorder.

mpegfilepath
The OS/2 directory path and filename of the MPEG-2 file. This file will be created or replaced when the next audiovisual program transmission is received by *localLterm*. If the file name is not supplied, then the default of D:\$TEMP$.M2T will be used.

Usage

- The avsprogram record command may only be issued against a logical terminal whose decoding device is of the type, filerecorder.
- When this command is issued, the logical terminal is opened and its devices are reserved for the terminal′s use. The terminal remains open and the devices reserved until an avsprogram stop command is issued for the logical terminal. Until that time, the terminal′s devices are unavailable for use by any other logical terminal.
- If the logical terminal is already recording an audiovisual program, the command will fail.
- If any of the logical terminal′s devices are currently being used by another logical terminal, the command will fail.

6.4.6 The avsprogram resume Command

This command resumes the playback of a recorded audiovisual program which was previously paused by using the **avsprogram pause** command.

Format

```
►►──avsprogram resume──sourceLterm──┬──────────────────┬──►◄
                                    └─-pos mmm:ss──────┘
```

Parameters

sourceLterm
The name of a local logical terminal that is playing the recorded audiovisual program. This terminal must have an encoding device of the type, filerecorder.

filepath
The OS/2 directory path and filename of the MPEG-2 file. This file will be created or replaced when the next audiovisual program transmission is received by *localLterm*.

-pos *mmm***:***ss*
The position from which to resume playing the program, expressed in terms of minutes (*mmm*) and seconds (*ss*). AVS translates this time into a byte offset into the MPEG-2 file and then begins playing the program from that point.

Usage

- The avsprogram resume command may only be issued against a program whose source is a locally-stored MPEG-2 file. The logical terminal that plays such a program uses an encoding device of the type fileplayer. Furthermore, the program being resumed must have been started by using the avsprogram start command, and paused by using the avsprogram pause command.
- Since the translation of minutes and seconds into a byte offset is only an estimation of the program′s play time, the playback position specified on the -pos parameter may not match the exact playback time, but it will fall within a few seconds of that time.
- When using the -pos causes AVS to change to a new byte offset in the MPEG-2 file, AVS does not attempt to locate the nearest I-Frame (complete image) to that offset. Therefore, the playback may resume in the middle of incomplete image information. Because of this, the video and audio may appear garbled for a moment while the receiving terminal regains synchronization with the MPEG-2 stream.

6.4.7 The avsprogram start Command

This command initiates the transmission of an audiovisual program from a local logical terminal to another audiovisual terminal (either local or remote).

Format

```
►►──avsprogram start──sourceLterm──┬─┤ destList ├─┬──┬─────────────┬──┬─────────────┬──────────►
                                   └─-b──destTerm─┘  └─-gop nn:p:b─┘  └─-pdu nnnnn─┘

►──┬────────────────┬──┬─────────────────────┬──┬──────────────────────────────────────┬──►◄
   └─-videobr nn.nn─┘  │            ┌─256─┐  │  └─-file mpegfilepath──┬────────────────┬─┘
                       └─-audiobr───┼─192─┼──┘                        └─-pos mmm:ss─┘
                                    └─384─┘
```

destList:

```
     ┌──────────────────────────┐
     ▼  ┌─-auto─┐     ┌───────┐ │
├────┴──┼───────┼─────▼destTerm┴─┴────┤
        ├─-h310─┤
        └─-vod──┘
```

Parameters

sourceLterm
The name of the logical terminal from which this program is to be transmitted.

destList
This is a list of destination terminal names to which a unidirectional program will be transmitted and optionally, information about how to connect to them. One or more such lists may be specified. Each list is grouped by the connection type, which is indicated by the connection type qualifier that precedes the list of terminal names. The first destination list may be specified with or without a preceding connection type qualifier. If no qualifier is specified on the first list, the -auto qualifier is assumed. All remaining lists must begin with a connection type qualifier. Multiple lists with similar connection type parameters are allowed.

The first terminal listed will be the first to receive the audiovisual program. The order of reception for the remaining terminals on the list is not predictable.

At most, one one of these may be another logical terminal on the local VaN.

-auto Indicates that the connection type is to be determined automatically. For each terminal in the sublist, AVS will first attempt to connect to the terminal using H.310 protocols. If the connection fails, AVS will then attempt to connect in ATM Forum Video-on-Demand (VoD) mode. If this connection type is used, the user does not have to know the exact type of each destination terminal. However, connecting in this mode causes a significant increase in the overall connection time for the addparty command.

If no connection qualifier is specified on the first destination sublist, -auto is assumed.

-h310 Tells AVS to connect to the terminals listed between this parameter and the end of the command line or a later -vod parameter using H.310 protocols. Use this option to reduce the processing and connection time for the addparty command when adding other VaNs to the transmission. Note, however, that use of this parameter requires you to know the connection type of the destination terminals.

-vod Tells AVS to connect to the terminals listed between this parameter and the end of the command line or a later -h310 parameter using ATM Forum Video-on-Demand (VoD) protocols. Use this option to reduce the processing and connection time for the addparty command when adding IBM MediaSTREAMERs, IBM Video Distribution Modules, or other ATM Forum VoD terminals to a transmission. Note, however, that use of this parameter requires you to know the connection type of the destination terminals.

destTerm Names of one or more audiovisual terminals to be added as receivers of the current program.

All of the following are valid destination lists:

```
ArlingtonHS QuincyHS AltonHS
ArlingtonHS -vod QuincyHS -h310 AltonHS
- vod ArlingtonHS QuincyHS AltonHS
- vod ArlingtonHS -h310 QuincyHS -vod AltonHS
- h310 ArlingtonHS -auto QuincyHS AltonHS
```

1. The first example defaults to –auto connection for all listed terminals.
2. The second example defaults to –auto connection for ArlingtonHS, uses a VoD connection to QuincyHS, and an H.310 connection to AltonHS.
3. The third example shows that the same connection qualifier may be used in multiple sublists. Here, both ArlingtonHS and QuincyHS use VoD connections.
4. The fourth example shows the use of the -auto qualifier.

-b *destTerm*
Specifies a single destination terminal with which a bidirectional (2-way) program will be established. This option is valid only if the local logical terminal is a receive-and-send terminal (RAST). Only one destination terminal (which also must be a RAST) may be listed. -b implies an H.310 connection type.

-gop *nn***:***p***:***b*
specifies the group of pictures configuration:

nn The number of frames per group

p The number of P frames to be included per I frame

b The number of B frames to be included per P frame

If the encoding device does not support B frames, then this portion of the specification is ignored.

-pdu *nnnnn*
Specifies the ATM packetized data unit (PDU) size in bytes.

The valid range is 384 to 65535. The PDU size should be an integral multiple of 188 (The size of a transport stream) + 8 bytes of overhead for the most efficient ATM transmission . The formula for this can be expressed as:

```
pdu size=(2+12n)*188+8
```

where n is any integer. For n=0, pdu=384. For n=1, pdu=2640, n2= pdu=4896, etc. The default is 8192.

-videobr *nn.nn*
Specifies the video bit rate in megabits per second (Mbps). The valid range is 4 to 15, although product implementation limitations significantly reduce the upper limit of this range. The default is 8.

-audiobr [**192 | 256 | 384**]
Specifies the audio bit rate for the transmission in kilobits per second (Kbps). The default is 256.

-file *mpegfilename*
The OS/2 directory path and filename of the MPEG-2 file. This file will be created or replaced when the next audiovisual program transmission is received by *localLterm*. If the file name is not supplied, then the default of D:\$TEMP$.M2T will be used.

-pos *mmm***:***ss*
The position from which to resume playing the program, expressed in terms of minutes (*mmm*) and seconds (*ss*). AVS translates this time into a byte offset into the MPEG-2 file and then begins playing the program from that point.

Usage

1. At most, one receiver of an outbound audiovisual program may be another logical terminal on the local VaN.
2. The addparty command will be attempted for every listed destination terminal. If errors occur for one or more of those terminals, an appropriate error message will appear and processing will continue on to the next terminal in the list.
3. Since the translation of minutes and seconds into a byte offset is only an estimation of the program′s play time, the playback position specified on the -pos parameter may not match the exact playback time, but it will fall within a few seconds of that time.
4. Using the **-pos** causes AVS to change to a new byte offset in the MPEG-2 file, AVS does not attempt to locate the nearest I-Frame (complete image) to that offset. Therefore, the playback may resume in the middle of incomplete image information. Because of this, the video and audio may appear garbled for a moment while the receiving terminal regains synchronization with the MPEG-2 stream.
5. Video quality for playback of locally recorded files through fileplayer devices is only guaranteed when played to a local destination terminal (that is, another logical terminal on the same VaN). While file playback will execute across an ATM connection, the quality of the playback is not guaranteed.

Examples:

The following command successfully starts a program between a local logical terminal named *OakRidge* and a remote terminal named *ETSUMainCampus* Since a connection type is not listed, AVS will first attempt to establish an H.310 connection and then, if necessary, a VoD connection.

```
[C:\]avsprogram start oakridge etsumaincampus
[C:\]
```

The next command successfully starts a program between a local logical terminal and three remote terminals, two using an H.310 connection and the last using a VoD connection. Additionally, the video bit rate is specified for low bandwidth consumption.

```
[C:\]avsprogram start etsumaincampus -h310 oakridge kingsport -vod greenville -videobr 6.3
[C:\]
```

The next command successfully starts a bidirectional program. Since a bidirectional program is indicated, an H.310 connection is used.

```
[C:\]avsprogram start -b kingsport oakridge
[C:\]
```

This command successfully plays a locally stored video clip named D:\announcment.m2t to a local decoder.

```
[C:\]avsprogram start clipplayer localtv -file d:\announcement.m2t
[C:\]
```

6.4.8 The avsprogram stop Command

This command terminates the audiovisual program currently being transmitted from or received by the specified logical terminal. Issue this command on either the transmitting or the receiving VaN.

Format

```
►►──avsprogram stop──localLterm──►◄
```

Parameters

localLterm
The name of a local logical terminal that is transmitting or receiving the audiovisual program to be stopped.

Usage

- In order to stop the program, it must have been started by using **avsprogram start**.
- If this command is issued to stop a program that is the source (root) of a multicast, transmission to all receiving terminals is terminated.
- If the program is being terminated by the last or only receiver, the program transmission also is terminated at the source.

6.4.9 The avsquery Command

This command lets you examine the status and configuration of your VaN. You also can examine current audiovisual transmissions and audiovisual device settings.

Format

```
►►──avsquery──┬─device──parameters─┬──►◄
              ├─lterm──parameters──┤
              └─van────────────────┘
```

Parameters

device
Displays configuration and status information for one or more audiovisual devices.

lterm
Displays configuration and status information for one or more logical terminals.

van
Displays general configuration information for the VaN.

Usage: Each form of the avsquery command is described below.

For all of the avsquery command examples, the following AVS configuration file will be assumed:

```
#
# Physical audiovisual devices....
device encoderOne   encoder      0
device decoderOne   decoder      0
device decoderTwo   decoder      0 backup
#
# Clip save/record devices.........
device recordFile   filerecorder
device playFile     fileplayer
#
# Virtual devices we wrote.........
device garysDevice  decoder virtual c:\ibmavs\vdevs\garydev.exe
#
# Logical terminals................
lterm  Room5Camera     15 encoderOne
lterm  Classroom5      16 decoderOne
lterm  Room5Conference 17 encoderOne decoderOne
lterm  ClipRecorder    21 recordFile
lterm  ClipPlayer      22 playFile
```

6.4.10 The avsquery Command

Displays configuration and status information for one or more audiovisual devices.

Format

```
                        ┌──────────────┐
                        ▼              │
►►──avsquery device──┴──────────────┴──┬────────┬──►◄
                        └─deviceName─┘     └─-stats─┘
```

Parameters

deviceName
The list of device names to be queried. These names must be defined on device records in the AVS configuration file. If no names are specified, all of the defined devices will be queried.

-stats
Indicates that detailed statistics for the devices are also to be displayed.

Examples:

The following command does a general query of all currently-defined audiovisual devices:

```
[C:\]avsquery device
Audiovisual Device: encoderOne
Type: IBM Encoder Adapter
Status: active
Current owner: Room5Camera
Card number: 0

Audiovisual Device: decoderOne
Type: IBM Decoder Adapter
Status: active
Current owner: Classroom5
Card number: 0

Audiovisual Device: decoderTwo
Type: IBM Decoder Adapter
Status: inactive
Current owner: -none-
Card number: 1
Other attributes: Backup device

Audiovisual Device: recoderOne
Type: File Recorder
Status: inactive
Not currently owned by a logical terminal

Audiovisual Device: garysDevice
Type: Virtual Decoder
Status: inactive
Current owner: -none-
Executable file: c:\ibmavs\vdevs\garydev.exe
[C:\]
```

The following command does a detailed query of the *encoderOne* device:

```
[C:\]avsquery device encoderOne -stats
Audiovisual Device: encoderOne
Type: IBM Encoder Adapter
Status: active
Currently owned by logical terminal: Room5Camera
Card number: 0
------ Stats ------
Bytes processed: 9587860488
Video format: (M)NTSC 3.58
Video bit rate (Mbps): 12.00
Audio bit rate (Kbps): 384
PDU size (bytes): 21848
[C:\]
```

6.4.11 The avsquery lterm Command

Displays configuration and status information for one or more logical terminals. If no logical terminal names are listed on the command, then information is displayed for all logical terminals currently defined on the VaN.

Format

```
              ┌──────────┐
              ▼          │
►►──avsquery lterm──┬──────────┬──┬────────┬──►◄
                    └─ltermName─┘  └─-stats─┘
```

Parameters

ltermName
The list of logical terminals to be queried. These names must be defined on the lterm records in the AVS configuration file. If no names are specified, all of the defined logical terminals will be queried.

-stats
Indicates that detailed statistics for the devices are also to be displayed.

Examples:

The following command does a general query of all currently-defined logical terminals:

```
[C:\]avsquery lterm

Logical Terminal: Room5Camera
Type: Send-Only
Status: Active
ATM Selector: 15
Devices: encoderOne

Logical Terminal: Classroom5
Type: Receive-Only
Status: Open
ATM Selector: 16
Devices: decoderOne

Logical Terminal: Room5Conference
Type: Receive-and-Send
Status: Inactive
ATM Selector: 17
Devices: encoderOne
       : decoderOne

Logical Terminal: ClipRecorder
Type: Receive-Only
Status: Inactive
ATM Selector: 21
Devices: recordFile

Logical Terminal: ClipPlayer
Type: Send-Only
Status: Inactive
ATM Selector: 22
Devices: playFile
[C:\]
```

The following command does a detailed query of the *room5camera* logical terminal:

```
[C:\]avsquery lterm room5camera -stats

Logical Terminal: Room5Camera
Type: Send-Only
Status: Active
ATM Selector: 15
Devices: encoderOne
------ Stats ------
Terminal connections: 1
Connection list:
   Classroom3
Transmission receiver list:
   Classroom3
[C:\]
```

6.4.12 The avsquery van Command

Displays general VaN configuration and operational parameters.

Format

```
►►──avsquery van──►◄
```

Examples:

```
[C:\]avsquery van
ATM address:  39.99.00.00.00.8B.00.00.00.21.00.3A.00.00.00.00.00.09
IP address: 122.16.57.231 (ralph.kramden.com)
TCP port base: 10000
Logical terminals: 5
Devices: 6
Supervision mode: Unsupervised
Terminal name/address mapping exit: Not registered
Maximum number of H245 sessions: 20
H245 timeout interval: 12 seconds
Active configuration file: C:\IBMAVS\ETC\CONFIG.VAN
ATM address resolution: C:\IBMAVS\ETC\AVSTERM.MAP
[C:\]
```

6.4.13 The avstermmap Command

This command lets you refresh the ATM terminal address resolution map. The filepath parameter specifies the OS/2 file path of the terminal name-to-address mapping file.

Format

```
►►──avstermmap──mapfilepath──►◄
```

Parameters

mapfilepath
The OS/2 directory and file name of the file that contains the mapping between audiovisual terminal names and ATM addresses.

Usage

- This command is intended to allow addition of new terminal names and addresses to an active audiovisual network. It is *not* intended for dynamically changing addresses of active terminals. Be careful to ensure that you do not modify the ATM addresses of audiovisual terminals that are currently involved in audiovisual program transmissions. If an active terminal's address is changed, the results will be unpredictable
- If a termNameToAddr mapping exit is registered, this command is processed but the exit still is used to resolve audiovisual terminal names and addresses.

6.4.14 The avsstart Command

This command starts up the Audiovisual Services software on the VaN.

Format

```
►►──avsstart──┬──────┬──┬─────────────┬──┬──────────────┬──┬──────────────┬──┬──────────────┬──►◄
              └─-fg──┘  └─-localtest──┘  └─-msg filepath─┘  └─-cfg filepath─┘  └─-map filepath─┘
```

Parameters

-fg This optional parameter specifies that the AVS process is to run in the foreground. The default is to run in the background.

-localtest
This optional parameter indicates that AVS is to be run in local test mode. In this mode, no ATM connections are used or even required. Audiovisual transmission may only be started between logical terminals within the local VaN. The -localtest option is intended for use as a verification tool when installing the VaN, as a diagnostic aid, and as a means of testing new AVS applications without affecting the rest of the audiovisual network.

-msg
This optional parameter specifies the file path to the audiovisual terminal directory file. The default is AVSTERM.MSG.

filepath
The OS/2 directory and file name of the indicated file.

-cfg
This optional parameter specifies the file path of the VaN configuration file. The default is AVS.CFG.

-map
This optional parameter specifies the file path of the terminal name to ATM address mapping file. The default is AVSTERM.MAP.

6.4.15 The avsstop Command

This command terminates the Audiovisual Services software.

Format

```
►►──avsstop──┬──────────┬──►◄
             └──-force──┘
```

Parameters

-force
Indicates that AVS is to terminate regardless of current activity. The default is to prompt the user if any logical terminals have active audiovisual programs or terminal connections.

Examples

```
avsstop

At least one terminal connection or audiovisual program is active.
Terminate anyway? (Y/N)

y

Audiovisual Services are shutting down.
```

6.4.16 The avstrace Command

This command activates, deactivates, or modifies AVS tracing.

Format

```
►►──AVSTRACE──┬──────────────────────────┬──┬────────────────────────────┬──┬────────┬──►◄
              │        ┌───────────┐     │  │          ┌───────────┐     │  └─-clear─┘
              └─-comp ─▼─component─┴─────┘  └─-events ─▼─tracepoint─┴───┘
```

Parameters

-comp *component*
This optional parameter indicates the list of components in which to activate the trace point. The component variable can be any combination of:

comm	communications subsystem
device	Device subsystem
terminal	Terminal manager
h245	H.245 protocol subsystem
api	API and callback subsystem
snmp	SNMP subsystem
services	System services subsystem and the configuration manager

-events *tracepoint*
This optional parameter indicates which event to track within each listed component. The tracepoint variable can be any combination of:

syscall	Operating system and service library calls
storage	Memory allocation and deallocation requests
call	Internal module entry and exit
comm	Communication operations (ATM, TCP/IP, and OS/2 named pipe calls)

-clear
This optional parameter indicates that the current trace settings are to be cleared before activating the new trace set according to the remaining parameters.

Chapter 7. Developing AVS Applications

This chapter describes the AVS programming environment and shows how to install the AVS Development Kit.

7.1 The AVS Programming Environment

AVS programs may be developed in the following environments:

- OS/2 Version 4
- Any C or C++ compiler that supports OS/2 Version 4 (for example, IBM Visual Age C++ Version 3.0).
- OS/2 Version 4 Software Developer′s Toolkit
- AVS Application Programming Support (comes preinstalled on each VaN and is a selectable component in the AVS installation program). This includes:
  - The required AVS C language header files
  - The required AVS import libaries
  - Programming examples, including two virtual devices that make it possible to run a subset of AVS on a stand-alone OS/2 machine for application development and testing.
  - A version of the AVS.CFG file for use in application development and testing

While the AVS software and AVS applications are intended to run only on a VaN platform, a subset of AVS may be run on a stand-alone OS/2 Version 4 machine for application development and test purposes. Specifically:

- Start AVS using the -localtest option so that no ATM connectivity is required. Additionally, use the special AVS configuration file named AVSDEV.CFG (located in the \IBMAVS\ETC directory) to provide a small, but usable AVS execution environment. The syntax of the command to do all of this is:

  ```
  avsstart -localtest -cfg c:\ibmavs\etc\avsdev.cfg
  ```

- The AVSDEV.CFG file defines two virtual devices named dummyenc (for dummy encoder) and dummydec (for dummy decoder). Each of these devices is implemented in a sample programs in \IBMAVS\SAMPLES:

  avsvenc.c Is a virtual encoder which generates fake transport stream data. Since the data is not a real transport stream, the output from this virtual device should not be transmitted to a real decoding device. However, the data is perfectly suitable for consumption by the dummydec decoder.

  avsvdec.c Is a virtual decoder which simply throws away all of the transport stream data it is given by AVS. Since the data is never used or verified, the fake transport streams generated by avsvenc.c are consumed as if it were a real transport stream.

  Using these two device definitions, two logical terminals are then defined: TESTSOT, which is a send-only terminal, and TESTROT, which is receive-only.

  Given the two logical terminals with their dummy devices, an AVS application may open either or both of the terminals, connect them in either H.310 or VoD mode, and start and stop audiovisual transmissions between the two. Additionally, you may use the the definitions in the supplied AVSDEV.CFG file as a basis for extending the test VaN environment as needed using multiple instances of the dummy virtual devices.

- Once your application programs run reliably in this test environment on a stand-alone OS/2 PC, you should then move them to a real VaN to test scenarios that require the full hardware environment or features that cannot be emulated on a single PC (such as multicast, for example).

7.2 Installing the AVS Development Kit

Get the five AVS diskettes and:

- Insert disk 1 into the floppy drive of your OS/2 V4 PC. and type `A:\INSTALL`.
- Once the installation utility comes up, the third screen (or so) will ask you to choose which components to install. Choose **Audiovisual Services** and **Application Programming Support**.

This will install the product tree, including the header files, sample programs, libraries, and executables that you will be able to use to test your code before you actually receive VaN boxes.

A couple of notes:

- After you install, but before you run AVS, you will need to add to IP host names to whatever resolution process you use. If AVS can′t resolve these IP host names, you′ll get error messages in the AVS server window to the effect that the terminals couldn′t be configured, and AVS will eventually terminate because it couldn′t define any logical terminals.

  We′ve found that the best way to set up the IP name resolution for these development test terminal names is:

  – Go into TCP/IP configuration and enable loopback support. The default loopback address it uses is 127.0.0.1.

  – While still in TCP/IP configuration, go to page 2 of the Host name panel, and click on **Look through HOSTS list before going to nameserver** option.

  – Create/update your hosts file with the two names TESTSOT and TESTROT. Each of these should use the loopback address. The entries are:

    ```
    127.0.0.1      TESTSOT
    127.0.0.1      TESTROT
    ```

    Note that OS/2 uses the ETC environment variable to identify the TCP/IP ETC path, so the easiest way of locating the hosts file is by invoking your editor with the path %etc%\hosts.

  If you want to, it′s fine to add these names to your domain name server, but if you do, you′re tying the names TESTSOT and TESTROT to one particular machine. If you have more than one person doing development (or are just using more than one machine for development/testing), this can get in the way.

- Once your TCP/IP configuration is reset and you′ve rebooted, use the following procedure to get your test environment going:

  1. Open the IBM VAN Audiovisual Services folder and double-click on the **AVS Local Test** icon. When run in local test mode, AVS ignores all of the

ATM stuff, so you'll be able to run without an ATM adapter or being active on an ATM network. You'll see a window pop up, several error messages appear, and finally, you should see the following message:

```
Audiovisual Services available and ready.
```

2. Open two other windows. In each, CD to \ibmavs\samples. In one, run avsvenc and in the other run avsvdec. These are two virtual devices, one a fake encoder and the other a fake decoder. In each window, you'll see messages like:

```
AVSVENC: A sample AVS virtual device (a virtual encoder)
application program
(C) Copyright International Business Machines Corp. 1997

Virtual device name: dummyenc
Callback function received NULL indication
Received rc = 0 (00000000) from avsRegister
Received rc = 0 (00000000) from avsOpenVirtualDevice
Waiting for VDEV START request
```

3. Now, from yet another session (!), you can start running your test code. Use TESTSOT and TESTROT as the logical terminal names on your tests.

To verify that the setup works after installation, you can issue the following commands from that third window:

```
avsprogram start testsot testrot -vod
```

When the prompt appears after this command, you should see that messages like the following have appeared in each of the virtual device windows:

```
VDEV QUERY STREAMS request received.
VDEV START request received.  Starting to receive transport stream data
==> 2 threads actively writing TS data...
```

To stop the transmission, enter:

```
avsprogram stop testsot
```

As a result, you should see something like the following in these device windows:

```
Unexpected rc = -15981824 (ff0c2300) from avsWriteTransportStream on thread 3
Thread 3 terminating: 55265232 TS bytes written.
Unexpected rc = -15981824 (ff0c2300) from avsWriteTransportStream on thread 4
Thread 4 terminating: 55215600 TS bytes written.
==> Total TS bytes written: 110480832
Waiting for VDEV START request
```

4. To shut down AVS, just hit Cntrl-C in the AVS server window, choose **Close** this session from the task list, You can use any of the other usual ways to knock down an OS/2 session. Note, however, that if you use Cntrl-C or Cntrl-break, the window will hang around until you explicitly close it.

Hopefully, these instructions will get you most of the way there.

Appendix A. VaN and VDM Technical Specifications

Table 13 (Page 1 of 7). Encoding

Description	Yes/No	Comments
COMPRESSION STANDARDS		
MPEG Standards		
MPEG-1 ISO/IEC 11172-1	No	MPEG-1 Program Stream is not supported.
MPEG-2 ISO/IEC 13818-1	Yes	Transport Stream format only.
MPEG-1 (ISO/IEC 11172-2)	Yes	MPEG-1 Video is supported MPEG-2 Transport Stream format only.
MPEG-2 (ISO/IEC 13818-2)	Yes	Simple Profile @ Main Level.
Video Conferencing Standards		
H.310	Yes	Supports all terminal (Receive and Send, Send- and Receive-Only terminal) for AAL-5 on Private Networks. Supports the H.245 control protocol as specified under H.310. Does not support any required and optional modes in H.310 that H.261 video encoding and G.711 audio encoding.
H.320	No	
VIDEO ENCODING ALGORITHM		
MPEG-2 Main Profile @ Main Level	Yes	Release 2
MPEG-2 Simple Profile @ Main Level	Yes	I-frame only, Release 1 I-P Frame, Release 2
MPEG-1 Stereo Audio	Yes	Single stereo channel, Release 1 Stereo channel plus two mono channels, Release 2
MPEG-2 or Dolby AC-3 Audio	No	
VIDEO ENCODING BIT RATE		
1.5 Mbps to 20 Mbps (indicate maximum rate if less than 20 Mbps) * I,P,B-frames * I-frame only	Yes Yes	Release 1 Release 2 Maximum Video Compressed Rate of 15 Mbps
Bit rate Encoding * Constant Bit rate * Variable Bit rate * Dynamic adjustment of bit rate while encoding * Automatic scene change I-frame insertion * Adaptive field/frame motion estimation	 Yes No No No Yes	
Capability for insertion of Picture Start Sequence Code	Yes	Programmable per number of frames between codes or on every I-frame.
Capability for Variable Buffer Verifier (VBV)	Yes	Default value is all F′s per encoder chipset.
Group of Pictures structure (GOP) and ratios must be user adjustable.	Yes	
VIDEO RESOLUTIONS		

Table 13 (Page 2 of 7). Encoding

Description	Yes/No	Comments
MPEG-1 SIF NTSC 352h x 240v x 30 fps	Yes	NTSC-M, NTSC-N, NTSC-4.43
MPEG-2 HALF CCIR-601 NTSC 352h x 480v x 30 fps	Yes	NTSC-M, NTSC-N, NTSC-4.43
MPEG-2 FULL CCIR-601 NTSC 704h x 480v x 30 fps	Yes	NTSC-M, NTSC-N, NTSC-4.43
MPEG-1 SIF PAL 352h x 288v x 25 fps	Yes	PAL-BGHI, PAL-N, PAL-M
MPEG-2 HALF CCIR-601 PAL 352h x 576v x 25 fps	Yes	PAL-BGHI, PAL-N, PAL-M
MPEG-2 FULL CCIR-601 PAL 704h x 576v x 25 fps	Yes	PAL-BGHI, PAL-N, PAL-M
MPEG-1 SIF FILM 352h x 240v x 24 fps	Yes	Via support of SMPTE 259M
MPEG-2 HALF CCIR-601 FILM 352h x 480v x 24 fps	Yes	Via support of SMPTE 259M, Release 2
MPEG-2 FULL CCIR-601 FILM 704h x 480v x 24 fps	Yes	Via support of SMPTE 259M, Release 2
MULTIPLEXING STREAM TYPES	Yes	
VIDEO ELEMENTARY STREAM	Yes	
AUDIO ELEMENTARY STREAM	No	
VIDEO and AUDIO ELEMENTARY STREAMS	No	
MPEG-2 PACKETIZED ELEMENTARY STREAM	Yes	
AUDIO PACKETIZED ELEMENTARY STREAM	Yes	
PID ASSIGNMENTS		
USER ASSIGNABLE	Yes	Follows conventions of ATSC A/53
MUX RATE		
1.5 Mbps to 20 Mbps	Yes	Maximum mux rate (Audio + Video + User) = 18 Mbps
PACKETS PER PACK		
USER ADJUSTABLE	No	Program Stream format not supported.
PACK SIZE		
USER ADJUSTABLE	No	Program Stream format not supported.
INTERMEDIATE CHROMA and LUMINANCE RESOLUTIONS		
4:2:0	Yes	
4:2:2	Yes	
4:4:4	No	
GRAPHICS, STILL FRAME IMAGES, TEXT		
CD QUALITY or BETTER	No	
AUDIO COMPRESSION STANDARDS		
ISO/IEC 11172 (PART 3) LAYER 2	Yes	
AUDIO Sampling Rate		
32 kHz	Yes	
44.1 kHz	Yes	

Table 13 (Page 3 of 7). Encoding		
Description	**Yes/No**	**Comments**
48 kHz	Yes	
AUDIO INPUT FORMATS		
TWO CHANNEL ANALOG	Yes	
TWO CHANNEL DIGITAL AES/EBU	Yes	Release 2
FOUR CHANNEL AUDIO	Yes	Release 2
STEREO ANALOG AUDIO	Yes	
JOINT STEREO AUDIO	Yes	
BALANCED AUDIO	Yes	Release 2. Can be supported in Release 1 product through the use of an external transformer.
UNBALANCED AUDIO	Yes	
VIDEO INPUT FORMATS		
SERIAL COMPONENT DIGITAL	Yes	SMPTE 259M, Release 2
SERIAL COMPOSITE DIGITAL	Yes	SMPTE 259M, Release 2
COMPONENT ANALOG	No	
COMPOSITE ANALOG	Yes	SMPTE 170M
SERIAL D1 (SMPTE 259M)	Yes	SMPTE 259M, Release 2
S-VIDEO	Yes	SMPTE 170M
EXTERNAL OUTPUT FORMATS		
Video Formats		
MPEG-1 System Streams	No	
MPEG-2 Program Streams	No	
MPEG-2 Transport Streams	Yes	Maximum mixed bit rate of 18 Mbps
Communication Transport Formats		
MPEG-2 Compressed Video packetized into ATM 155 Mbps OC-3c AAL-5 layer cells * Permanent Virtual Connection (PVC) Setup * Switched Virtual Connection (SVC) Setup * Point-to-Point * Point-to-Multipoint * Multi-mode Fiber * Single-mode Fiber * UNI 3.1 or higher Signaling * Adjustable PDU Size	Yes	Supports SVC, Pt-Pt, Pt-Mpt Multi-mode Fiber and UNI 3.1 with adjustable PDU Size (4K-64K).
MPEG-2 Compressed Video packetized into the following Communication Transport Protocols: * High Speed Serial Interface (HSSI) * ISDN (64 - 1920 Kbps) * T-1 (1.544 Mbps) * T-3/DS3 (44.736 Mbps) * E-1 (2.048 Mbps) * E-3 (34.368 Mbps) * ATM 25 Mbps * ATM OC-1 51.840 Mbps * ATM 100 Mbps * ATM OC-3 155 Mbps	Yes	Supports ATM 25 Mbps, ATM DS-3, ATM OC-3 interfaces.

Table 13 (Page 4 of 7). Encoding		
Description	**Yes/No**	**Comments**
EXTERNAL COMMUNICATION INTERFACES FOR EXTERNAL OPERATIONAL, SYSTEMS and NETWORK MANAGEMENT CONTROL		
ETHERNET IEEE 10baseT	Yes	
REMOTE DIAL-IN via MODEM (9600 or higher)	Yes	
In-Band ATM using low-level ATM APIs	Yes	Supports RFC 1577 (Classic IP over ATM) for in-band ATM communications for external, operational and systems control.
EXTERNAL COMMUNICATION PROTOCOL		
TCP/IP	Yes	Supports TCP/IP through both the LAN interface and the ATM interface (through Classic IP).
MPEG-2 PRIVATE DATA AREA		
Vertical Blanking Interval information for Lines 11-21 (e.g., closed-captioning in Line 21) must be preserved during encoding into the Private Data Area for later re-insertion by the decoding operation.	Yes	Closed-captioning/EDS supported (EIA-608) Data services on lines 10-18 supported (EIA-516).
OPERATIONS MANAGEMENT		
Graphical User Interface (GUI) for manual initiation of encoding operations.	Yes	Includes an Audiovisual Services API that permits the development of any application on the platform. A GUI will be provided to manually initiate encoding application on the platform. This software is provided as a sample application to demonstrate the capabilities of the platform as well. The source code of the application will be provided to allow customization of the application.
An application program interface (API) for initiating encoding operations from an external control system.	Yes	Is supplied with the Audiovisual Services Applications Programming interface. This API provides access to all H.310 related services and extensions to this standard to permit acknowledged audiovisual session management (See copy of System Level Design Document for specifications of the API).
Remote login or operations from an external control system via a LAN using TCP/IP protocols.	Yes	Supports TCP/IP. Thus, an operator may log in remotely using telnet. Upon login, the available applications are the Software Installer Operations command line interface.
System/Network Management		
Support for SNMP MIB network management protocols. * Configuration Management * Alarms and Fault Management * System Status * Detection of loss of service * Password Security	Yes	Provides an SNMP MIB-II for network management purposes.
External Monitoring of Encoding Operation		

Table 13 (Page 5 of 7). Encoding		
Description	**Yes/No**	**Comments**
Monitoring support for the source input and the real-time encoding output.	Yes	Encoding source can be viewed in analog format by monitoring video loopthru connector.
Diagnostics		
All diagnostic and error alerts/traps must be stored and available for retrieval both locally and remotely to a network management system.	Yes	All diagnostic error and traps will be available in a logfile for lo cal and remote retrieval. In addition, the traps will be sent to the manager of the system through SNMP.
Logging and audit trail information must be stored and available for retrieval both locally and remotely.	Yes	All logging information is stored and is available for retrieval locally and remotely. The remote retrieval may be performed using ftp. The IBM Video Access Node does not contain an audit trail for billing purposes.
Dynamic tracing which can be stored and available for retrieval both locally and remotely is a preferred option.	Yes	All tracing information is stored and is available for retrieval locally and remotely. The remote retrieval may be performed using ftp.
Video Filters		
External preprocessing or filtering?	No	
ENVIRONMENTAL REQUIREMENTS		
Redundant and highly-available packaging. Describe features and how they are invoked * N+1 Encoding Adapters in the system * Re-routing around a failed component * Power Supplies * Cooling Fans * Hot-Pluggable Components	 No No Yes Yes Yes	Platform redundancy features dependent on platform selection.
Mountable in a standard 19-inch rack-mount enclosure.	Yes	Height = 5 RU (rack units)
Front and/or rear access	Yes	Rear access standard, front access available via patch panel.
External cabling or patch panel.	No	
Ability to operate with the following accessories attached: * Display Monitor * Keyboard * Mouse * Display/Keyboard/Mouse Switch if multiple systems installed within the same rack	Yes	
Optional Components: * External Control or Management System * Attachment for a tape storage unit. Describe in detail including type, capacity and attachment interface.	Yes	Permits the use of an external management system to manage a network of encoders. The specific manager is not provided.
Local disk storage, if yes, please describe in detail including quantity, capacity, attachment interface, redundancy.	Yes	2 GB standard. platform provides four 3.5″ 1″ height bays and one 5.25″ half-height bays for additional hard files, all supported by Ultra F/W SCSI.

Table 13 (Page 6 of 7). Encoding		
Description	**Yes/No**	**Comments**
statistical multiplexing	No	Does not perform any statistical multiplexing of individual video streams into the ATM network. However, will support pseudo-variable bit rate operation. This will allow statistical multiplexing at the network level.
Encryption/Decryption	No	Does not support encryption or decryption of audiovisual data.
Autoranging Power Supply * 110-240 Vac, 50-60 Hz * 2-meter detachable plug	No	A manual switch is provided to select proper voltage range of the AC input voltage. Input frequency is 50 ± 3 Hz or 60 ± 3 Hz. MINIMUM RANGE: 90-180 VAC MAXIMUM RANGE:137-265 VAC
Front Panel Indicators * Power on * Power off * Encoding Operation * Fault or Error	 Yes Yes No Yes	The IBM PC Server front panel indicators are as follows: LED1= Power ON LED2= Unattended LED LED3= POST Active LED4= SCSI HF Active LED5= CPU1 Active LED6= CPU2 Active LED7= ENET Xmit LED8= ENET Rcv LED9= ENET Link LED10= Security
Front Panel Controls * Locking Power Off/On Switch	No	C2 Security is provided on the IBM PC Server: Mechanical lock on the system cover. Separate power-on and keyboard password Unattended Start Mode allows system keyboard to be locked, but at the same time allow network computer access System Management Adapter for Reporting Trouble (SMART) card provides the following additional features: Remote power on/off of system unit Monitoring of system temperature in two areas Monitoring of system fan rotation Ability to drive system speaker Ability to turn on security LED to indicate tampering has occurred.
Required Connectors align. * Composite (BNC) * Digital Serial Component (BNC) * Component (BNC, R-Y/B-Y/Y) * S-VIDEO (Mini-DIN)	 Yes Yes No Yes	align. Release 1 Release 2 Release 1
Temperature * Operational: 5°C to 35°C * Storage: -20°C to 80°C	No	 * Operational: 10°C to 35°C * Storage: 1°C to 60°C * Shipment: -40°C to 60°C

Table 13 (Page 7 of 7). Encoding		
Description	**Yes/No**	**Comments**
Humidity * Operational - 10% to 80% non-condensing * Storage - 10% to 80% non-condensing	 Yes Yes	* Operational - 8% to 80% * Storage - 5% to 80% (Including condensation but excluding rain)
Electrical and Safety Standards and Approvals * Meet or exceed North American UL Standards * Meet or exceed North American FCC Part 15 Class A * Meet or exceed North American CSA * Meet or exceed North American EMI * Applicable support for international standards at a future date	Yes	
NEBS Compliance	No	Optional platform available if required.
-48V Operation	No	Optional external equipment support for -48VDC to 110 VAC conversion.
Built-in Battery/UPS Backup	No	Optional external equipment support for UPS if required.

Table 14 (Page 1 of 8). Decoding		
Description	**Yes/No**	**Comments**
COMPRESSION STANDARDS		
MPEG Standards		
MPEG-1 ISO/IEC 11172-1	Yes No	IBM Video Distribution Module does not support MPEG-1 System Stream format.
MPEG-2 ISO/IEC 13818-1	Yes No	IBM Video Distribution Module does not support MPEG-2 Program Stream format.
MPEG-1 (ISO/IEC 11172-2)	Yes Yes	Must be in Transport Stream wrapper for IBM Video Distribution Module.
MPEG-2 (ISO/IEC 13818-2)	Yes	
Video Conferencing STANDARDS		
H.310	Yes	IBM Video Access Node is same support as Encoder. IBM Video Distribution Module does not support H.245.
H.320	No	
VIDEO DECODING BIT RATE		
1.5 Mbps to 20Mbps (indicate maximum rate if less than 20 Mbps)	Yes	Transport Stream format limited to 15 Mbps video.
Dynamic adjustment of decoding video rate to insure proper synchronization	Yes	IBM Video Access Node according to OpenMPEG standard. IBM Video Distribution Module performs dynamic adjustment of decoding process.
VIDEO RESOLUTIONS		
MPEG-1 SIF NTSC 352h x 240v x 30 fps	Yes	
MPEG-2 HALF CCIR-601 NTSC 352h x 480v x 30 fps	Yes	

Table 14 (Page 2 of 8). Decoding		
Description	**Yes/No**	**Comments**
MPEG-2 FULL CCIR-601 NTSC 704h x 480v x 30 fps	Yes	
MPEG-1 SIF PAL 352h x 288v x 25 fps	Yes	
MPEG-2 HALF CCIR-601 PAL 352h x 576v x 25 fps	Yes	
MPEG-2 FULL CCIR-601 PAL 704h x 576v x 25 fps	Yes	
MPEG-1 SIF FILM 352h x 240v x 24 fps	No	
MPEG-2 HALF CCIR-601 FILM 352h x 480v x 24 fps	No	
MPEG-2 FULL CCIR-601 FILM 704h x 480v x 24 fps	No	
GRAPHICS, STILL FRAME IMAGES, TEXT		
Overlay still images, video and text of any size/shape at any display location.	Yes	Hardware capabilities exist for this function. Software API support will be developed pending more specific requirements.
Display MPEG still images.	Yes	Hardware capabilities exist for this function. Software API support will be developed pending more specific requirements.
Text - English and European Fonts.	Yes	Hardware capabilities exist for this function. Software API support will be developed pending more specific requirements.
GUI Interface for Text Overlay Formatting	Yes	Hardware capabilities exist for this function. Software API support will be developed pending more specific requirements.
AUDIO COMPRESSION STANDARDS		
ISO/IEC 11172 (PART 3) LAYER 2	Yes	
AUDIO Sampling Rate		
32 kHz	Yes	
44.1 kHz	Yes	
48 kHz	Yes	
AUDIO OUTPUT FORMATS		
TWO CHANNEL ANALOG	Yes	
TWO CHANNEL DIGITAL AES/EBU	No	
FOUR CHANNEL AUDIO	Yes	Video on channel B same as channel A. Reduces available ports on IBM Video Distribution Module by 1/2.
STEREO ANALOG AUDIO	Yes	
JOINT STEREO AUDIO	Yes	
BALANCED AUDIO	Yes	
UNBALANCED AUDIO	Yes	
VIDEO INPUT FORMATS		

Table 14 (Page 3 of 8). Decoding

Description	Yes/No	Comments
MPEG-2 Compressed Video packetized into ATM 155 Mbps OC-3c AAL-5 layer cells * Permanent Virtual Connection (PVC) Setup * Switched Virtual Connection (SVC) Setup * Point-to-Point * Point to Multipoint * Multi-mode Fiber * Single-mode Fiber * UNI 3.1 or higher Signaling * MPEG-1 System Streams * MPEG-2 Packetized Elementary Streams * MPEG-2 Program Streams * MPEG-2 Transport Streams	Yes	IBM Video Access Node: same as Encoder IBM Video Distribution Module: No support for MPEG-1 system streams, MPEG-2 PES, or MPEG-2 Program Streams.
Compressed Video packetized into the following Communication Transport Protocols: * High Speed Serial Interface (HSSI) * T-1 (1.544 Mbps) * T-3/DS3 (44.736 Mbps) * E-1 (2.048 Mbps) * E-3 (34.368 Mbps) * ATM 25 Mbps * ATM OC-1 51.840 Mbps * ATM 100 Mbps * ATM OC-3 155 Mbps	Yes	IBM Video Access Node: same as Encoder IBM Video Distribution Module: all interfaces supported by the IBM 8260/8285.
ANALOG INPUT for PICTURE-IN-PICTURE OVERLAY	No	
EXTERNAL VIDEO OUTPUT FORMATS		
UNCOMPRESSED * ANALOG NTSC * ANALOG PAL * ANALOG AUDIO * D1 DIGITAL	Yes Yes Yes No	
EXTERNAL COMMUNICATION INTERFACES FOR EXTERNAL OPERATIONAL, SYSTEMS and NETWORK MANAGEMENT CONTROL		
ETHERNET IEEE 10baseT	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported, see 8260 or 8285.
REMOTE DIAL-IN via MODEM (9600 or higher)	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported. See 8260 or 8285.
In-Band ATM using low-level ATM APIs	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported. See 8260 or 8285.
EXTERNAL COMMUNICATION PROTOCOL		

Table 14 (Page 4 of 8). Decoding		
Description	**Yes/No**	**Comments**
TCP/IP	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported. See 8260 or 8285.
MPEG-2 PRIVATE DATA AREA		
Data stored in the MPEG-2 private data area in the encoded input stream must be extracted during the decoding operation for re-insertion into the correct Vertical Blanking Interval (lines 11-21). An example is closed-captioning information on Line 21.	Yes	IBM Video Access Node: Closed-captioning/EDS (EIA-608), Release 1 IBM Video Distribution Module: Closed-captioning/EDS (EIA-608), Release 1 IBM Video Access Node: Data services on lines 10-18 (EIA-516), Release 2 IBM Video Distribution Module: Data services on lines 10-18 (EIA-516), Release 2
OPERATIONS MANAGEMENT		
Graphical User Interface (GUI) for manual initiation of decoding operations	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Once ATM connection is made, port is ready to receive and decode stream, no "start" required.
An application program interface (API) for initiating decoding operations from an external control system.	Yes No	IBM Video Access Node: Same as Encoder Video Distribution Module: Low-level ATM connections performed by server/codec.
Remote login or operations from an external control system via a LAN using TCP/IP protocols.	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported. See 8260 or 8285.
System/Network Management		
Support for SNMP MIB network management protocols? List and describe what types of information can be queried, configured or changed. * Configuration Management * Alarms and Fault Management * System Status * Detection of loss of service * Password Security	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported.
External Monitoring of Decoding Operation		
Monitoring support for the source input and the real-time decoding output.	Yes	Source input is monitored via quality measurements on ATM interface. Decoder can provide quality monitoring of decompression process. No on-card monitoring of analog signal quality is provided.
Diagnostics		

Table 14 (Page 5 of 8). Decoding		
Description	**Yes/No**	**Comments**
All diagnostic and error alerts/traps must be stored and available for retrieval both locally and remotely to a network management system.	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Local retrieval only. Remote supported by 8260.
Logging and audit trail information must be stored and available for retrieval both locally and remotely.	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Local retrieval only. Remote supported by 8260.
Dynamic tracing which can be stored and available for retrieval both locally and remotely is a preferred option.	Yes No	IBM Video Access Node: Same as Encoder IBM Video Distribution Module: Not supported.
ENVIRONMENTAL REQUIREMENTS		
Redundant and highly-available packaging. Describe features and how they are invoked. * N+1 Decoding Adapters in the system * Re-routing around a failed component * Power Supplies * Cooling Fans * Hot-Pluggable Components	 Yes No Yes Yes Yes	IBM Video Access Node: Provision spare ports as backup. Platform redundancy features dependent on platform selection. IBM Video Distribution Module: Provision spare ports as backup. Otherwise, all features available on 8260 and 8285 platforms.
Mountable in a standard 19-inch rack-mount enclosure	Yes	
Encryption/Decryption	No	
Autoranging Power Supply * 110-240 Vac @ 50-60 Hz * 2-meter detachable plug	VaN: No VDM: Yes	IBM Video Access Node: A manual switch is provided to select proper voltage range of the AC input voltage. Input frequency is 50±3 Hz or 60±3 Hz. MINIMUM RANGE: 90-180 VAC MAXIMUM RANGE:137-265 VAC

Table 14 (Page 6 of 8). Decoding

Description	Yes/No	Comments
Front Panel Indicators * Power on * Power off * Encoding Operation * Fault or Error	 Yes Yes No Yes	IBM Video Access Node: The IBM PC Server front panel indicators are as follows: LED1= Power ON LED2= Unattended LED LED3= POST Active LED4= SCSI HF Active LED5= CPU1 Active LED6= CPU2 Active LED7= ENET Xmit LED8= ENET Rcv LED9= ENET Link LED10= Security IBM Video Distribution Module: Status lights for port activity, power, and blade status are provided. The 8260/8285 also provide diagnostic display capabilities.
Front Panel Controls * Locking Power Off/On Switch	No	
Front and/or rear access	No Yes	IBM Video Access Node: Rear access standard. IBM Video Distribution Module: Front access only.
External cabling or patch panel	Yes	IBM Video Access Node: Front access provided via available 2 RU patch panel.
Ability to operate with the following accessories attached * Display Monitor * Keyboard * Mouse * Display/Keyboard/Mouse Switch if multiple systems installed within the same rack		
IBM Video Access Node IBM Video Distribution Module: Not applicable to the 8260/8285 platform.	Yes	

Table 14 (Page 7 of 8). Decoding		
Description	**Yes/No**	**Comments**
IBM Video Access Node Optional Components * External Control or Management System	Yes	C2 Security is provided on the IBM PC Server: Mechanical lock on the system cover. Separate Power-on and Keyboard Password Unattended Start Mode allows system keyboard to be locked, but at the same time allow network computer access System Management Adapter for Reporting Trouble (SMART) card provides the following additional features: Remote power on/off of system unit Monitoring of system temperature in two areas Monitoring of system fan rotation Ability to drive system speaker Ability to turn on Security LED to indicate tampering has occurred.
IBM Video Distribution Module Optional Components * External Control or Management System	Yes	The 8260/8285 platform provides the following: C2 Security Intelligent power Intelligent cooling See 8260/8285 Documentation
IBM Video Access Node Required Connectors * Composite (BNC) * Digital Serial Component (BNC) * Component (BNC, R-Y/B-Y/Y) * S-VIDEO (Mini-DIN) * GENLOCK	 Yes Yes No Yes Yes	
IBM Video Distribution Module Required Connectors * Composite (BNC) * Digital Serial Component (BNC) * Component (BNC, R-Y/B-Y/Y) * S-VIDEO (Mini-DIN) * GENLOCK	 Yes No No No Yes	Release 1 Release 2
IBM Video Access Node Temperature * Operational: 5°C to 35°C * Storage: -20°C to 80°C	No	* Operational: 10°C to 35°C * Storage: 1°C to 60°C * Shipment: -40°C to 60°C

Table 14 (Page 8 of 8). Decoding

Description	Yes/No	Comments
IBM Video Distribution Module (8260) Temperature * Operational: 5°C to 35°C * Storage: -20°C to 80°C	No	* Operational: 10°C to 50°C * Storage: 1°C to 60°C * Shipment: -40°C to 60°C
IBM Video Access Node Humidity * Operational - 10% to 80% non-condensing * Storage - 10% to 80% non-condensing	Yes Yes	* Operational - 8% to 80% * Storage - 5% to 80% (Including condensation but excluding rain)
IBM Video Distribution Module (8260) Humidity * Operational - 10% to 80% non-condensing * Storage - 10% to 80% non-condensing	Yes Yes	* Operational - 8% to 95% * Storage - 5% to 80%
Electrical and Safety Standards and Approvals * Meet or exceed North American UL Standards * Meet or exceed North American FCC Part 15 Class A * Meet or exceed North American CSA * Meet or exceed North American EMI * Applicable support for international standards at a future date	Yes	IBM Video Access Node and IBM Video Distribution
NEBS Compliance	No	IBM Video Access Node: Optional platform available if required. IBM Video Distribution Module: Depth of platform is non-conformant.
-48V Operation	No Yes	IBM Video Access Node: Optional external equipment support for -48VDC to 110 VAC conversion. IBM Video Distribution Module: Supported by 8260/8285 platform
Built-in Battery/UPS Backup	No	Optional external equipment support for UPS if required.

Appendix B. Special Notices

This publication is intended to help Customers, IBM technical professionals, service specialists, marketing specialists, and marketing representative to demonstrate, install, administer and support the IBM 8300 Video Access Node (VaN) and the IBM 8260 Video Distribution Module (VDM). The information in this publication is not intended as the specification of any programming interfaces that are provided by 8300 VaN and 8600 VDM. See the PUBLICATIONS section of the IBM Programming Announcement for the 8300 VaN and 8600 VDM for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	IBM
MediaSTREAMER	NetView
Nways	Operating System/2
OS/2	RS/6000

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

AC-3	Dolby Laboratories Licensing Corporation
C + +	American Telephone and Telegraph Company, Incorporated
CSA	Canadian Standards Association
DCE	The Open Software Foundation
Dolby	Dolby Laboratories Licensing Corporation
Hayes	Hayes Microcomputer Products, Incorporated
NFS	Sun Microsystems Incorporated
RCA	General Electric Company

Appendix C. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 187.

- *Campus ATM Design Guidelines*, SG24-5002-00
- *Asynchronous Transfer Mode (ATM) Technical Overview*, SG24-4625-00
- *Understanding and Using the IBM MSS Server*, SG24-4915-00
- *Internetworking over ATM: An Introduction*, SG24-4699-00

C.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

C.3 Other Publications

These publications are also relevant as further information sources:

- *Image and Video Compression Standards Algorithms and Architectures*, by Vasudev Bhaskaran and Konstantinos Konstantinides
- *Graphics File Formats Reference and Guide*, by C. Wayne Brown and Barry J. Shepherd
- *MPEG-2 Transport over ATM Networks*, by Christos Tryfonas, University of California, Santa Cruz
- *Background Information on MPEG-1 & MPEG-2 Television Compression*, found at URL:

  http://www.cdrevolution.com/text/mpeginfo.htm

- *The MPEG standard*, found at URL:

  http://www.crs4.it/~luigi/MPEG

- *Video Access Node*, found at URL:

http://www.networking.ibm.com/van

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
(+45) 48 14 2207 (long distance charge)	Outside North America

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com
- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number

- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

Index

R

S

T

V

ITSO Redbook Evaluation

IBM Networked Video Solution Over ATM Implementation
SG24-4958-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redeval@vnet.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction ____________

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:
____

____

____

____

What other redbooks would you like to see published?
____

____

____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)
____

____

____

____

____

Printed in U.S.A.

SG24-4958-00