

WHITE PAPER

July 1998

Compaq Computer Corporation
ECG Emerging Markets and Advanced Technology

CONTENTS

VIDEO TECHNOLOGY	3
The Human Eye	3
Analog Video.....	4
Analog Composite Video.....	4
Analog Component Video	4
Digital Video	5
Digital Video Formats	5
Network Delivery Challenges	6
The Bandwidth Problem.....	6
Scaling.....	7
Compressing—Codecs.....	8
Video Codec Standards.....	9
H.261.....	9
H.263.....	9
JPEG and MJPEG.....	9
MPEG.....	10
VIDEO-STREAMING	12
Isochronous Video.....	13
Video Streaming System	14
Network Considerations.....	15
LANs/Intranets	17
Public Internet	18
Public Broadband Networks.....	20
VIDEO SERVERS	22
Application Software.....	22
Video Server Hardware.....	23
High Capacity Disk Storage.....	24
High Sustainable Throughput.....	24
High Performance Network.....	24
Multiple CPUs	24
Expandable System Memory.....	24
High Availability.....	24
Rack-mount for Easy Access.....	25
Attractive Cost per Stream.....	25
Example: Compaq ProLiant.....	25
ACRONYMS	27
RESOURCES	28

Video Streaming Technology

If a picture is worth a thousand words, then a video is worth a thousand pictures. The sights and sounds of video teach us, entertain us, and bring our fantasies to life. While text, graphics, and animation provide for interesting content, people naturally gravitate to the richer and more realistic experience of video. That is because video—with audio—adds the ultimate level of realism to human communication that people have come to expect from decades of watching moving pictures in the real-world media of TV and movies.

As all such real-world media continues to migrate toward "everything digital", video too is becoming digital. Video delivery has evolved from the analog videotape format of the 1980s to a digital format delivered via CD-ROM, DVD-ROM and computer networks. As a series of digital numbers, digital video has the advantage of not degrading from generation to generation, and, because it can reside on a computer disk, it is easy to store, search, and retrieve. It can also be edited and easily integrated with other media such as text, graphics, images, sound, music, as well as transmitted without any loss in quality. And now it is possible to deliver digital video over computer networks including corporate Intranets and public Internets directly to desktop computers.

What makes this network delivery possible is the emergence of new technology called "video-streaming". Video streaming takes advantage of new video and audio compression algorithms as well as new real-time network protocols that have been developed specifically for streaming multimedia. With video streaming, files can play as they are downloaded to the client, thus eliminating the necessity to completely download the file before playing, as has been the case in the past. This has the advantages of playing sooner, not occupying as much disk space, minimizing copyright concerns, and reducing the bandwidth requirements of the video.

This white paper discusses the salient characteristics of this new video-streaming technology. How the properties of human vision shape the requirements of the underlying video technology. How the high bit rate and high capacity storage needs of video drives the demand for high video compression and high bandwidth networks. How the real-time nature of video demands the utmost in I/O performance for high levels of sustained throughput. How the characteristics of video streaming shapes the requirements of video server hardware. And finally, how Compaq video streaming servers meet these demanding requirements through high performance I/O architectures and the adherence to industry standards.

Please direct comments regarding this communication to the ECG Emerging Markets and Advanced Technology Group at: EMAT@compaq.com

COMPAQ

NOTICE

THE INFORMATION IN THIS PUBLICATION IS SUBJECT TO CHANGE WITHOUT NOTICE AND IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Compaq, Deskpro, Compaq Insight Manager, ProLiant, Netelligent, and SmartStart are registered with the United States Patent and Trademark Office.

Microsoft, Windows, Windows NT, Windows NT Server 4.0, Terminal Server Edition, Microsoft Office 97, Microsoft Excel, Microsoft Word, Microsoft Outlook 97, Microsoft Internet Explorer, PivotTable, Microsoft Visual Basic, Microsoft SQL Server, Microsoft Exchange, and Internet Information Server are trademarks and/or registered trademarks of Microsoft Corporation.

Pentium, Pentium II, and Pentium Pro are registered trademarks of Intel Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1998 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Video Streaming Technology



First Edition (July 1998)

ECG068/0798

Video Technology

The Human Eye

An understanding of video technology starts with an understanding of the properties of the human eye. This is because the electronic eye of the video camera tries to mimic what the human eye sees. Basically, the human eye detects, or “sees”, electromagnetic energy in the visible light frequency spectrum ranging from a wavelength of about 400 nanometers (nm) to 700 nm. The eye cannot detect electromagnetic radiation outside this spectrum.




Visible Light

The human eye detects this light energy through the use of photoreceptors known as "rods" and "cones". There are approximately 120 million rods distributed across the spherical surface called the "retina" at the back of the eye. These rods are sensitive to light intensity. By contrast, there are only about 8 million cones that are sensitive to color on the surface of the retina. These cones tend to be centrally located at the back of the eye. The large number of light-sensing rods, compared to color-sensing cones, makes the eye much more sensitive to changes in brightness than to changes in color. This fact is taken advantage of in the video and image compression schemes discussed below which sample color at a lower rate than that of brightness. This is also why night vision, which relies on the low light intensity sensing capability of rods, is devoid of color, and why peripheral vision, which is not directed at the center of the retina is not as color sensitive.




Brightness changes detected by 120M rods



Color changes detected by 8M cones

For color sensing there are three types of cones capable of detecting visible light wavelengths. It has been determined that a minimum of three color components—e.g., Red, Green, and Blue—corresponding to the three types of cones, when properly filtered, can simulate the human sensation of color. Since color does not exist in nature—it is literally in the eye and brain of the beholder—these cones sense light in the visible spectrum and our brain processes the result to provide us with the sensation of color. This process is additive in that the brain can create colors—e.g., red + blue = purple—that don’t exist in the pure spectrum. These properties of human vision are used in video compression schemes, as well as in display systems to provide efficient methods for storing, transmitting, and displaying video data.

Another characteristic of human vision important to video technology is that of “image persistence”. This is where an image remains on the retina, even though the original object has been physically removed or replaced by another image. This persistence tends to be around .1 second. This causes the eye to perceive motion if the image is changing at a rate of greater than 10 frames per second (fps). It has been determined that smooth motion requires a frame rate > 15 fps.




Analog Video

Analog Composite Video


Analog video represents video information in frames consisting of fluctuating analog voltage values. In early analog video systems individual video signals— brightness, sync, and color— were all combined into one signal known as "composite" video. This composite signal can be transmitted over a single wire. Compared to other forms of video, composite analog video is lowest in quality. "Compositing" can result in color bleeding, low clarity and high generational loss when reproduced.

Analog Component Video

The low quality of 1-wire composite video gave way to higher quality "component" video where the signals are broken out into separate components. Two of the most popular component systems are the Y/C 2-wire system and the RGB 3-wire system.




The Y/C system separates the brightness or luminance (Y) information from the color, or chroma (C) information. This approach—called "S-Video"—is used in Hi-8 and Super VHS video cameras. The RGB system separates the signal into three components—Red, Green, and Blue—and is used in color CRT displays.



Another approach to component video is to use Luminance (Y), Hue (U), and Saturation (V) as the three components. Hue describes the color's shade or tone, and saturation the "purity" or "colorfulness" of the color.

This approach dates back to the introduction of color TV. For color TVs to be backward compatible and for black and white TVs to be able to

receive color signals, the color and brightness components were separated. Thus black and white TVs could subtract out the chroma—hue and saturation—information of a color signal and color TVs could display only the luma information received from a black and white transmission. This enabled both types of TVs to peacefully coexist. In turns out that YUV signals can be transformed into RGB signals and vice-versa by using simple formulas.



Digital Video

A major disadvantage of analog video is that it tends to degrade from one generation to the next when stored or reproduced. Another is that it often contains imperfections in the form of "artifacts" such as "snow" in the picture due to noise and interference effects. In contrast to analog video, digital video represents the video information as a series of digital numbers that can be stored and transmitted error free without degrading from one generation to the next. Digital video is generated by sampling and quantizing analog video signals. It may therefore be composite (D2 standard) or component (D1 standard) depending on the analog source. Until recently, digital video has been mostly stored on sequential tape because of the high capacity requirements, but advances in magnetic and optical disk capacity and speed make it economically feasible to store video on these media. To do this, analog video may be "captured" by digitizing it with a capture card and storing it as digital video on a PC's hard drive.

This makes it possible to more easily retrieve it, search through it, or edit it. Recently, new digital camcorders have emerged that store the video in digital form directly in the camcorder—usually on tape, but sometimes on a disk in the camcorder itself. Digital video from these sources may go directly to the hard drive of a PC by using an appropriate interface card.

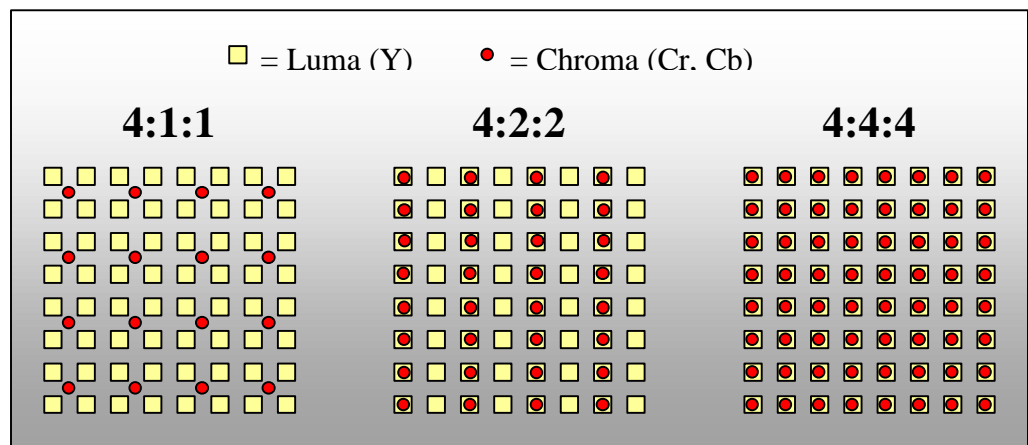
The quality of digital video may be judged based on three main factors:

1. **Frame Rate**—The number of still pictures displayed per second to give the viewer perception of motion. The National Television Standards Committee (NTSC) standard for full motion video is 30 frames per second (fps)—actually 29.97 fps—where each frame is made up of odd and even fields, hence 30 fps = 60 fields per second. By comparison, film is 24 fps
2. **Color Depth**—The number of bits per pixel for representing color information. For example, 24-bits can represent 16.7 million colors, 16-bits around 65,535 colors, or 8-bits only 256 colors.
3. **Frame Resolution**—Typically expressed as the width and height in pixels. For example, a full screen PC display is 640x480; a quarter screen is 320x240, a one-eighth or "thumbnail" is 160x120.

Digital Video Formats

To mimic the eye's perception of color, computer monitors display color information about each pixel on the screen using the RGB (Red, Green, Blue) format. Digital video, however, often uses a format known as YCrCb, where Y represents a pixel's brightness, or "luma", and Cr represents the color difference Red - Y, and Cb represents the color difference Blue - Y. By subtracting out the luminance Cr and Cb represent "pure" color. Together CrCb are referred to as "chroma".

Video Formats



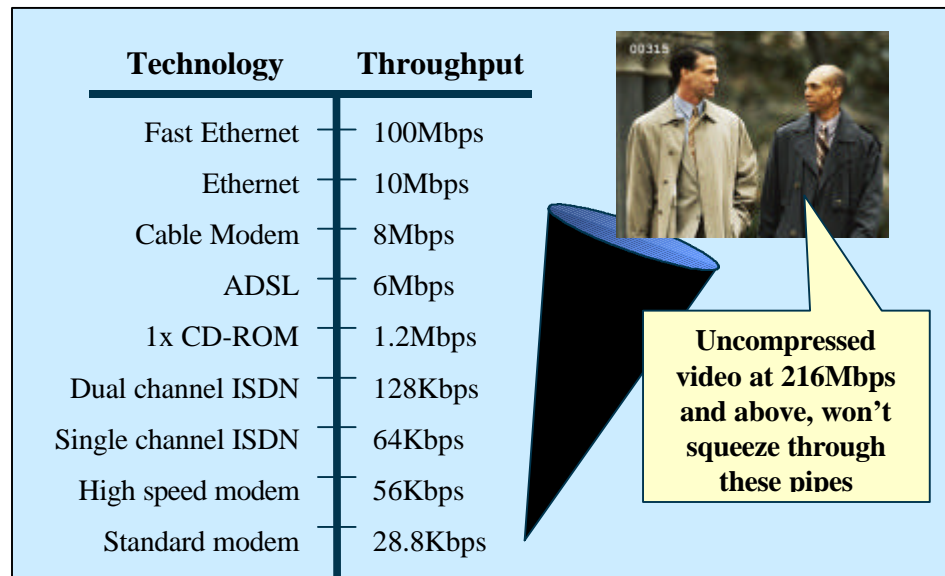
The '4' in the above descriptions indicates that luma is sampled at 4 times the basic 3.375 MHz frequency and the '1' and '2' indicates that the chroma is sub-sampled at 1 or 2 times the basic frequency. This approach to storing and transmitting video has the advantage of enabling file size reduction without any noticeable impact to the human eye on picture quality. Since, the eye detects subtle variations in brightness easier than differences in color, more bits are typically used to represent brightness with fewer bits representing color information. In this scheme each video pixel has its own luma, or Y, value, but groups of pixels may share CrCb chroma values. Even though some color information is lost, it is not noticeable to the human eye. Depending on the format used, this conversion can result in a 1/3 to 1/2 reduction in file size. The color bits required per pixel for each of these formats is 24 (4:4:4), 16 (4:2:2), and 12 (4:1:1).

Network Delivery Challenges

The bandwidth required by video is simply too great to squeeze through narrow data pipes. For example, full-screen/full-motion video can require a data rate of 216 MegaBits per second (Mbps). This far exceeds the highest data rate achievable through most networks or across the bus of older PCs. Until recently, the only practical ways to get video on a PC were to play video from a CD-ROM or to download a very large file across the network for playback at the user's desktop. Neither of these approaches is acceptable for delivery of content across a network.

The Bandwidth Problem

The scope of this problem can be seen by looking at the following illustration of available bandwidth for several methods of data delivery.



As can be seen from this illustration, even a high bandwidth Ethernet LAN connection cannot handle the bandwidth of raw uncompressed full-screen/full-motion video. A substantial amount of video data compression is necessary.

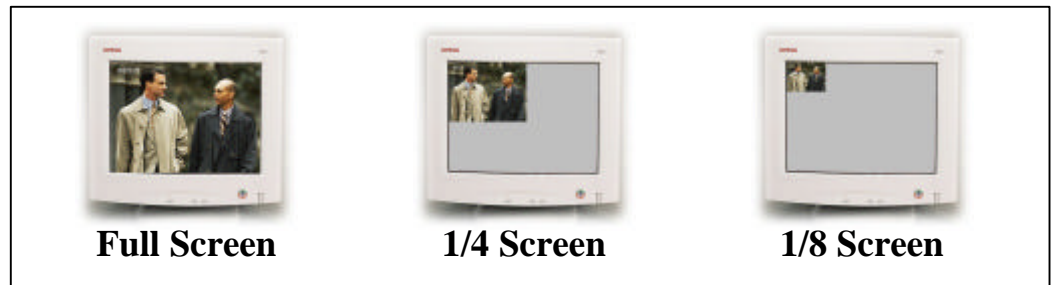
Successfully delivering digital video over networks can involve processing the video using three basic methods:

1. **Scaling** the video to smaller window sizes. This is especially important for low bandwidth access networks such as the Internet, where many clients have modem access.

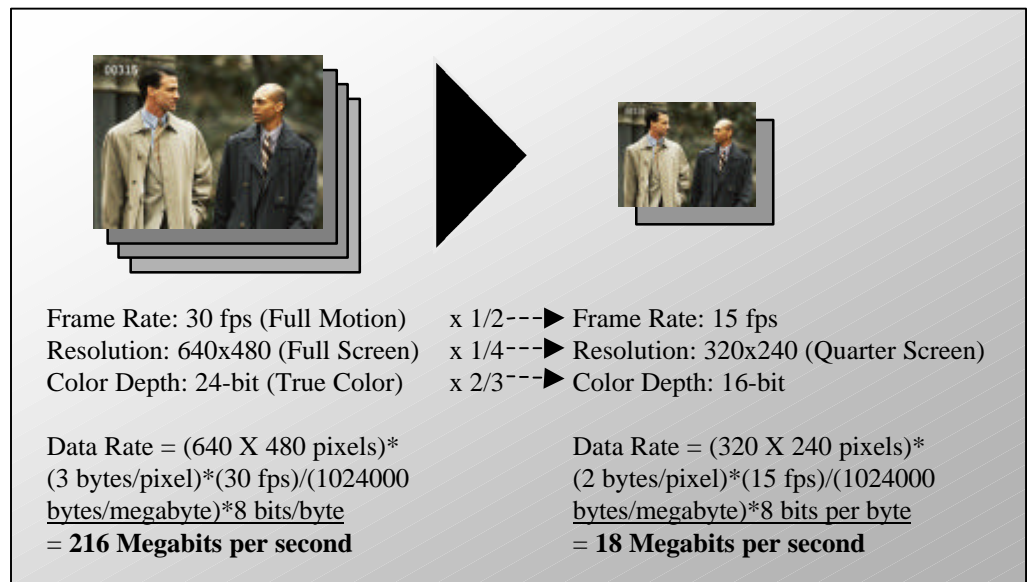
2. **Compressing** the video using lossy compression techniques. This is generally needed for almost all networks because of the high bandwidth requirements of uncompressed video.
3. **Streaming** the video using data packets over the network. Small video files may be downloaded and played, but there is a tendency to stream larger video content for faster viewing.

Scaling

While converting from the RGB color space to a subsampled YCrCb color space helps reduce file size, it is only a 1/3 to 1/2 reduction which is not nearly enough. Techniques to lower this further involve scaling one or more of the three factors mentioned above: frame rate, color depth, and frame resolution. For example, scaling the frame resolution results in different size windows for showing the video on the screen.



Further scaling of all three parameters can dramatically reduce the video rate as can be seen in the following diagram.

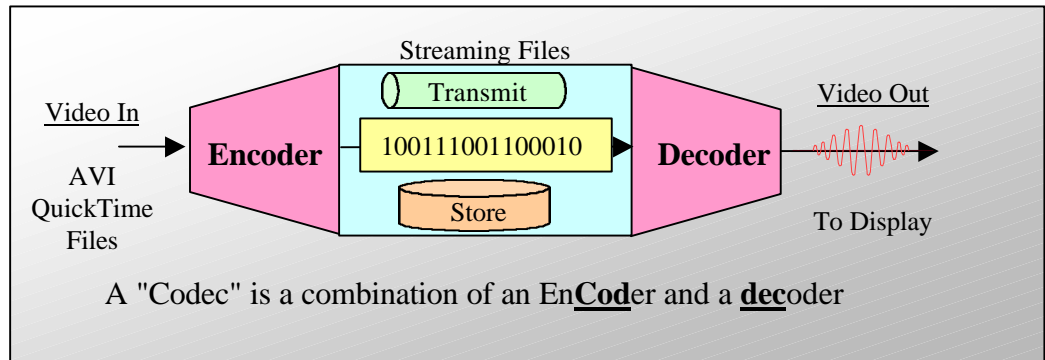


Even though the above scaling represents over a 10:1 reduction in data rate at the expense of size and video quality, it is still not enough for most network delivery. For example, a 10BaseT Ethernet network supports data rates of 10 MegaBits/sec. This is not enough bandwidth to deliver even one video stream at the above scaled data rate. Further scaling can be done. For example, the video can be scaled to a "thumbnail" size video at a few frames per second with 8-bit color—but this is poor in quality and still does not accomplish the data rate reduction necessary to

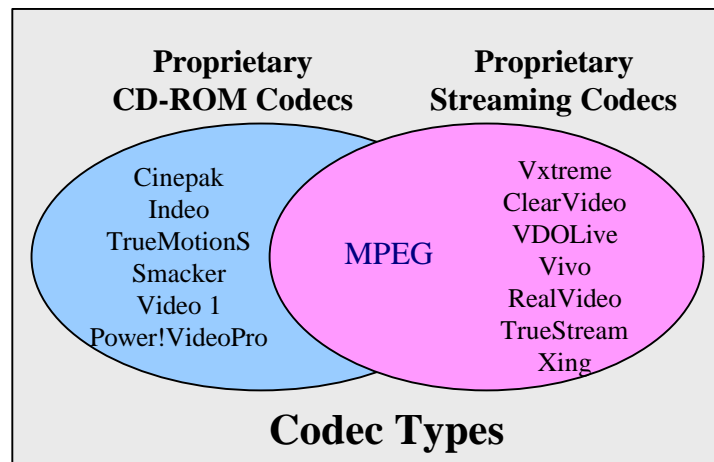
deliver very many streams for most network delivery. To achieve further reduction in data rate, video compression is needed.

Compressing—Codecs

Different algorithms and techniques known as "codecs" have been developed for compressing video signals. Video compression techniques take advantage of the fact that most information remains the same from frame to frame. For example, in a talking head video, most of the background scene typically remains the same while the facial expressions and other gestures change. Taking advantage of this enables the video information to be represented by a "key frame" with "delta" frames containing the changes between the frames. This is typically called "interframe" compression. In addition, individual frames may be compressed using lossy algorithms similar to JPEG photo-image compression. An example of this is the conversion from RGB to the Y/C color space described above where some color information is lost. This type of compression is referred to as "intraframe" compression. Combining these two techniques can result in up to 200:1 compression. This compression is achieved through the use of a "codec"—an encoder/decoder pair, depicted as follows.



Codecs vary depending on their purpose, for example, wide bandwidth vs. narrow bandwidth or CD-ROM vs. network streaming. Encoders generally accept file types such as Audio/Video Interleave (AVI) and convert them into proprietary streaming formats for storage or transmission to the decoder. Multiple files may be produced corresponding to the various bit rates supported by the codec. A codec may also be *asymmetric* or *symmetric* depending on whether it takes longer to encode than decode. Some codecs are very compute intensive on the encode side and are used primarily for creating content once that will be played many times. Symmetric codecs, on the other hand, are often used in real-time applications such as live broadcasts. A number of codecs have been developed specifically for CD-ROMs while others have been developed specifically for streaming video.



Video Codec Standards

H.261

The H.261 video-only codec standard was created by the ITU in 1990 for global video phone and video conferencing applications over ISDN. It was designed for low bit rates assuming limited motion as is typical with videophone applications. It was also assumed that ISDN would be deployed worldwide. Since each ISDN B Channel has a data rate of 64 Kbps, H.261 is also sometimes referred to as "Px64" where P can take integer values from 1 to 30. For compatibility between different TV systems—NTSC, PAL, SECAM—a Common Intermediate Format (CIF) was defined that will work across displays for all of these systems. CIF and Quarter-CIF resolution are defined as:

H.261 Resolutions	
Format	Resolution
QCIF:	176 x 144
CIF:	352 x 288

H.261 frame rates can be 7.5, 10, 15, or 30 fps.

H.261 has been the most widely implemented video conferencing standard in North America, Europe, and Japan, and formed the starting point for the development of the MPEG-1 standard described below.

H.263

H.263 was developed by the ITU in 1994 as an enhancement to H.261 for even lower bit rate applications. It is intended to support videophone applications using the newer generation of PSTN modems at 28.8 Kbps and above. It benefits from the experience gained on the MPEG-1 standard. It supports five picture formats:

H.263 Resolutions	
Format	Resolution
Sub-QCIF:	128 x 96
QCIF:	176 x 144
CIF:	352 x 288
4CIF:	704 x 576
16CIF:	1408 x 1152

Bitrates range from 8 Kbps to 1.5 Mbps. H.263 is the starting basis for MPEG-4 discussed below.

JPEG and MJPEG

JPEG stands for "Joint Photographic Experts Group." This group developed a compression standard for 24-bit "true-color" photographic images. JPEG works by first converting the image from an RGB format to a YCrCb format described above to reduce the file size to 1/3 or 1/2 of its original size. It then applies a sophisticated algorithm to 8x8 blocks of pixels to round off and quantize changes in luminance and color based on the properties of the human eye that detects subtle changes in luminance more than in color. This lossy compression technique has compression ratios in the range of 2-30.

MJPEG stands for "motion JPEG" and is simply a sequence of JPEG compressed still images to represent a moving picture. Video capture boards sometimes use MJPEG since it is an easily editable format, unlike MPEG, described next. A disadvantage of MJPEG is that it does not handle audio.

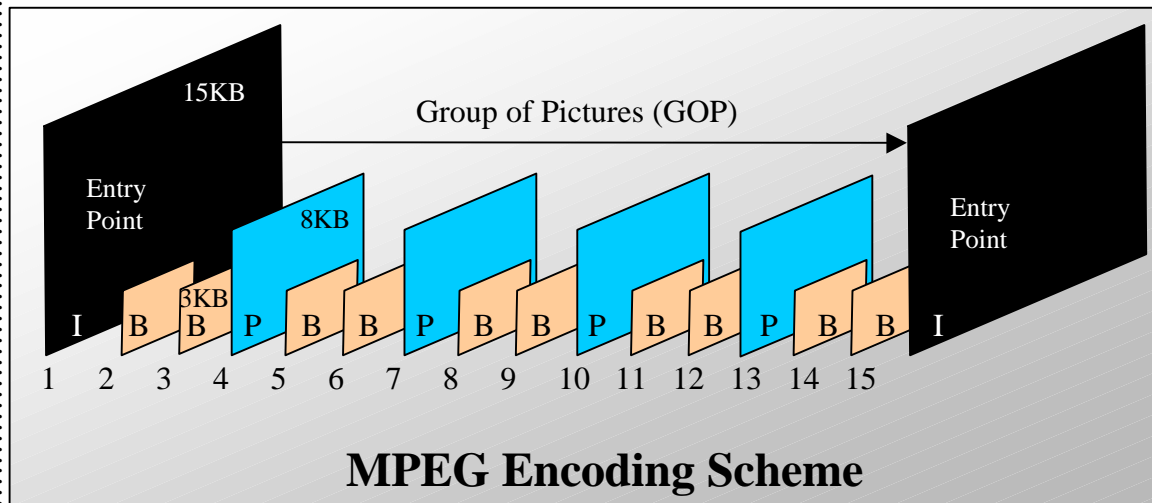
MPEG

The International Standards Organization (ISO) has adopted a series of video codec standards known as MPEG. MPEG stands for Moving Pictures Experts Group, a workgroup of the International Standards Organization (ISO). This group defined several levels of video compression standards known as MPEG-1, MPEG-2, and MPEG-4.

MPEG-1

The MPEG-1 standard was defined in January 1992 and was aimed primarily at video conferencing, videophones, computer games, and first generation CD-ROMs. MPEG-1 is also the basis of VideoCD and CD-i players. It was designed to provide consumer-quality video and CD-ROM quality audio at a data rate of approximately 1.5 Mbps and a frame rate of 30 fps. Originally, MPEG-1 was designed for playing video from 1x CD-ROMs and to be compatible with data rates of T1 lines.

MPEG-1 uses *Inter*frame compression to eliminate redundant information between frames, and *Intra*frame compression within an individual frame using a lossy compression technique similar to JPEG image compression. Three types of frames result from the compression algorithm: I-frames, P-frames, and B-frames. These are illustrated in the following diagram.



I-frames do not reference other previous or future frames. They are stand-alone as "Independent" frames and are compressed only with intraframe encoding. They are thus larger than the other frames. When video is played or indexed into with rewind and fast forward, they are also the entry points into the video since only they represent a complete picture.

P-frames, on the other hand contain "Predictive" information with respect to previous I or P frames, containing only the pixels that have changed since the I or P-frame, including taking motion into account. P-frames are therefore smaller than I-frames. I frames are sent at a regular interval, e.g., about every 400ms, and P-frames are sent at some time interval after the I-frame that varies based on the implementation.

Depending on the amount of motion in the video, P-frames may not come fast enough to give the perception of smooth motion. To compensate for this, B-frames are therefore inserted in between the I and P-frames. B-frames are "bi-directional" in that they use information in the previous I or P-frame as well as in the future I or P-frame. The information they contain is interpolated between two endpoints rather than representing actual data, making the assumption that the pixel information will not change drastically between the end-points. As a result these "dishonest" B-frames contain the most amount of compression and are the smallest frames. In order for a decoder to decode these B-frames it must have the corresponding I and P-frames they are based on, hence the frames may be transmitted out of order to reduce decoding delays.

A frame sequence consisting of an I-frame and its subsequent B and P-frames before the next I-frame—typically around 15 frames—is called a Group of Pictures (GOP). An I-frame of the GOP acts as the basic entry access point.

The result of this compression technique is that MPEG format is not easy to edit since you cannot enter the video at any point. Also the quality of the resulting video will depend on the particular implementation and the amount of motion in the video. The MPEG encoding/decoding process can also be seen as computationally intensive which requires the use of specialized hardware or a PC with a powerful processor.

The compression achieved using this technique enables about 72 minutes of video on a single CD-ROM—not quite enough for a full-length feature movie, which typically needs about 2 hours or two CD-ROMs. In addition to compression techniques, the MPEG-1 standard also supports playback functions such as fast forward, fast reverse, and random access into the bitstream. However, as mentioned above, these access points are only at the I-frame boundaries. The CD-ROM quality audio is stereo (2 channel) 16-bit sampled audio at 44KHz. The main MPEG-1 resolution of 360 x 242 is called Standard, or Source, Input Format (SIF) and, unlike the Common Intermediate Format (CIF) defined earlier, differs for NTSC and PAL systems. The computer industry has defined its own square pixel version of the SIF format at 320 x 240.

Originally, MPEG-1 decoding was done by hardware add-in boards with dedicated MPEG audio and video decoder chips because of the processing power required. Currently, MPEG-1 can be decoded in software on a Pentium 133 or better.

MPEG-2

MPEG-2 was adopted in the Spring of 1994 and is designed to be backward compatible with MPEG-1. It is not designed to replace MPEG-1, but to enhance it as a broadcast studio-quality standard for HDTV, cable television, and broadcast satellite transmission. Resolution is full-screen (ranging from NTSC 720x480 to HDTV 1280x720) playback with a scan rate of 60 fields per second. The latter enables MPEG-2 to support interlaced TV scanning systems as well as progressive scan computer monitors. Compression is similar to MPEG-1, but is slightly improved with the video exceeding SVHS quality, versus VHS for MPEG-1. Audio is also improved with six channel surround-sound versus CD-ROM 2-channel stereo audio for MPEG-1. The use of Interframe compression similar to MPEG-1 makes fast-motion scenes the most difficult to compress. Because of this, MPEG-2 supports two encoding schemes depending on the needs of the application: variable bit rate, to keep the video quality constant, and varying quality to keep the bit rate constant.

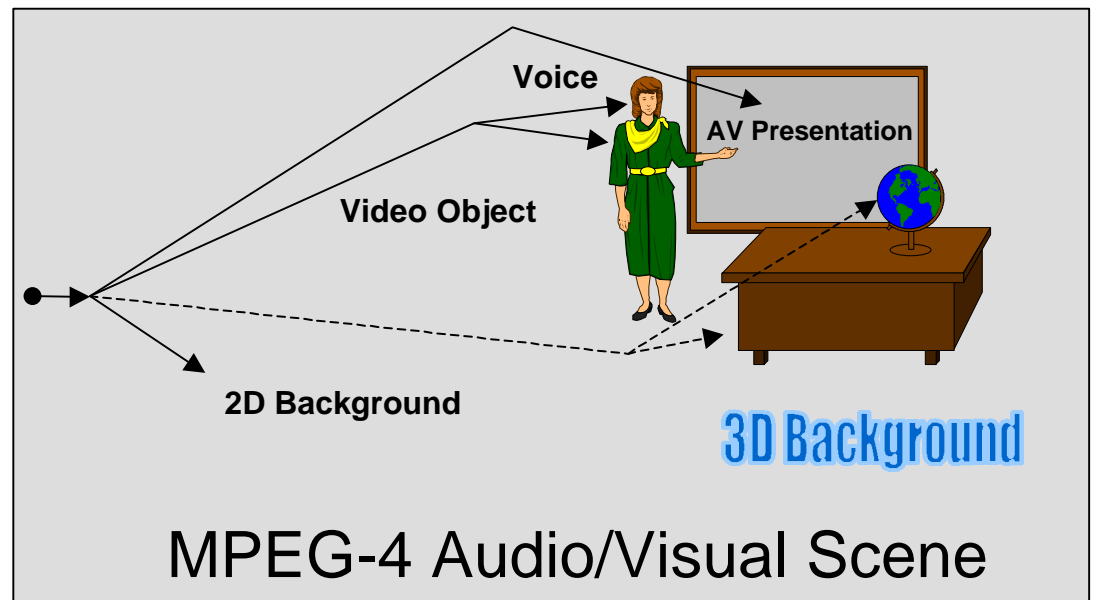
The variable data rates for MPEG-2 range from 2Mbps to 10Mbps. This typically requires a 4x CD-ROM drive for playback. A 4x CD-ROM can only store only 18 minutes of video versus a 1x CD with 72 minutes of video for MPEG-1. MPEG-2 was selected as the core compression

standard for DVD-ROM which, at 7x the capacity of the CD, is ultimately expected to replace CD-ROM and VHS tapes for long-play quality video. DVD-ROM can hold at least 130 minutes of MPEG-2 video with Dolby AC-3 surround-sound stereo audio. (DVD-ROM supports Dolby AC-3 audio rather than the MPEG-2 audio).

Another significant difference between MPEG-2 and MPEG-1 is the support of “Transport Streams” in MPEG-2. The concept of Transport Streams is designed for the delivery of video through error-prone networks. These streams consist of smaller fixed-length 188-byte data packets for error control and can also contain multiple video programs that are not necessarily time related. This makes MPEG-2 more suitable for the delivery of video over error-prone networks such as coax cable TV, ATM, and satellite transponders.

MPEG-4

MPEG-4 is a new MPEG standard that was originally proposed in 1993 for low-bit-rate applications such as the Internet or PSTN for such applications as video-conferencing, Internet video phones, video terminals, video email, wireless mobile video devices, and interactive home shopping. It was originally designed to support data rates of 64Kbps or less, but recently has been enhanced to support a wide range of bit rates from 8 Kbps to 35 Mbps. This enables it to support both consumer and professional video with a variety of resolutions corresponding to the H.263 video conferencing standard discussed earlier. Its main difference is the ability to transmit objects described by shape, texture, and motion, instead of just the transmission of rectangular frames of pixels. This enables interaction with multimedia objects as opposed to the frame-based interaction of MPEG-1 and MPEG-2.



This object-based streaming approach has application in configurable TV, interactive DVD, interactive web pages, and animations.

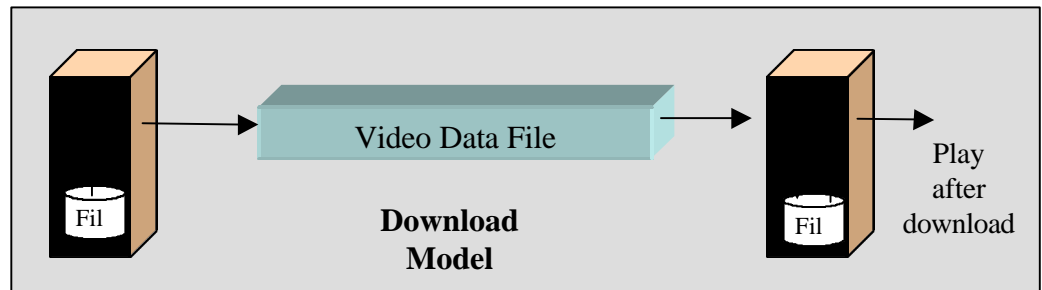
Video-Streaming

Until fairly recently, video has been delivered by the “download and play” method. In this approach, the entire video file is downloaded over a network to the client and stored on a hard drive. When it is all there, the user can play it from the hard drive. The advantage of this approach is that reasonably high quality video can be delivered, even over a low bandwidth


network. The disadvantage is that the user may have to wait for a long time for the download to occur, plus have a large amount of disk storage space available to store the video file. There are also copyright concerns since a copy of the video is made to the user's hard drive.

To address these issues, a new technology has recently been developed for delivering scaled and/or compressed digital video across a network called "video-streaming." Video streaming takes advantage of advances in video scaling and compression techniques as well as the use of network protocols that have been developed for real-time media streaming. These new protocols enable video with synchronized audio to begin playing at the client desktop before the entire file is received. This has the advantage of playing the video sooner and not requiring a large amount of disk space on the client. In addition, copyright concerns are reduced, since a complete copy does not have to reside on the user's hard drive. Another advantage is that video content may be integrated and streamed along with other media and the video can play inside standard Web browsers.

The following diagrams contrast the "download and play" model with the new "streaming" model.



Delivering video by the download & play method involves transmitting the complete file—e.g., in AVI format—to the client. These files can be very large and the user must wait until the complete file is downloaded before playing.



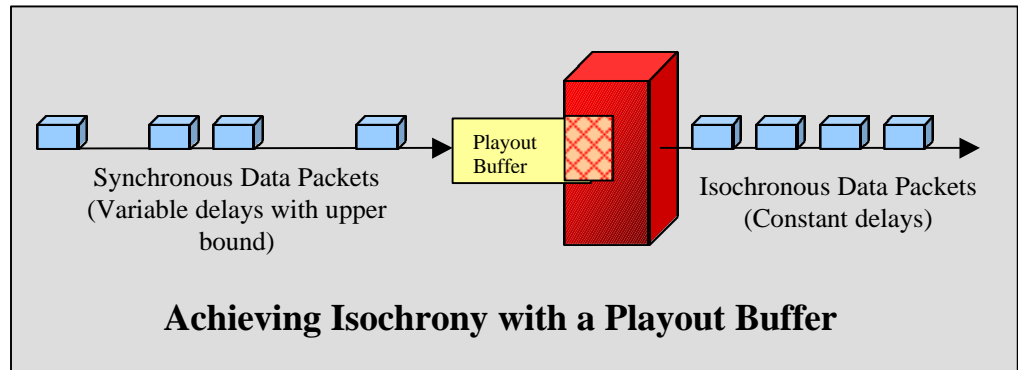
By contrast, with video streaming, a second or two of packetized video is buffered at the client, and then the video begins to immediately play out of the buffer as the file streams in. The user does not have to wait for the whole file to download before viewing it.

Isochronous Video

In order to play smoothly, video data needs to be available continuously and in the proper sequence without interruption. This is difficult to do over packet-switched networks such as the

Internet since individual packets may take different paths and arrive at different times at the destination. In some cases, packet errors or lost packets may also result in re-transmission causing disruption of the video. On Intranets such as shared Ethernet video packets may also have to contend with other data traffic such as file transfers, database access, access to shared modem and fax, etc. The continuous streaming data needs of video are difficult to meet in this type of environment.

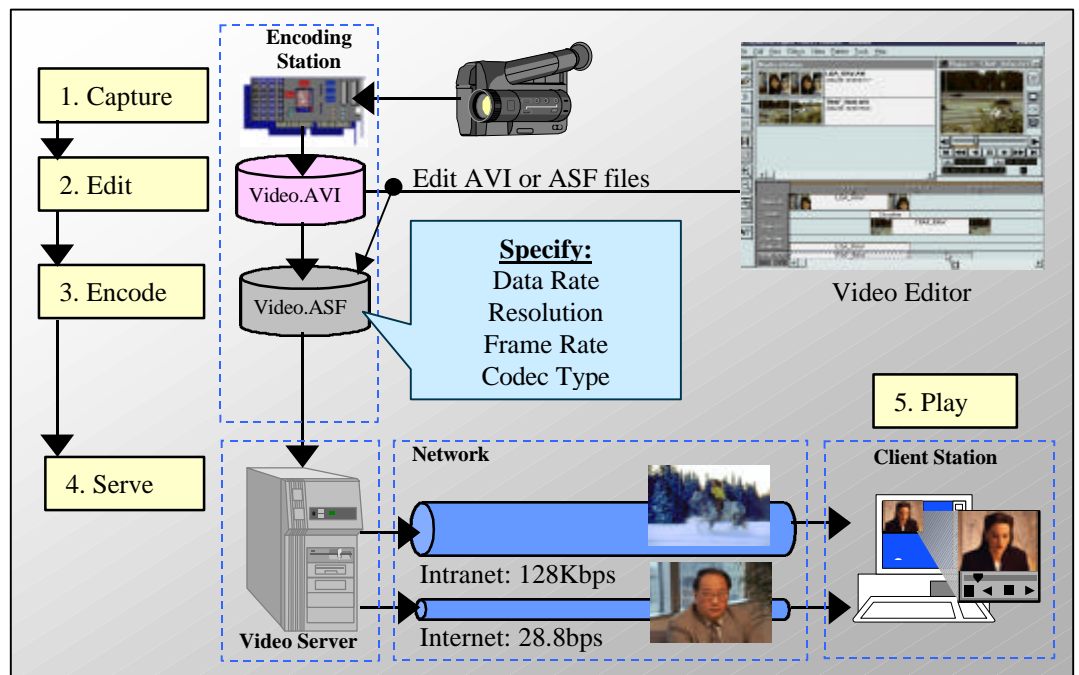
A network is called "isochronous" if it provides constant transmission delay and "synchronous" if it provides bounded transmission delay. For the latter, isochrony can be achieved with a "playout" buffer.



One way to achieve this effect is to insert time stamps and sequence numbers in the video data packets so they can be properly assembled and played out of the receiving buffer. Over IP-based networks such as the Internet, a new protocol called "Real-Time Protocol" (RTP) has been developed to accomplish this.

Video Streaming System

A complete video-streaming system involves all of the basic elements of creating, delivering, and ultimately playing the video content. The main components of a complete video streaming system to accomplish this—Encoding Station, Video Server, Network Infrastructure, and Playback Client—are illustrated in the following diagram.



Step 1. Capture: As this diagram shows, the first step in the process of creating streaming video is to "capture" the video from an analog source such as a camcorder or VHS tape, digitize it and store it to disk. This is usually accomplished with an add-in analog video capture card and the appropriate capture software. Newer digital video sources such as digital video camcorders can be captured straight to disk with a "Firewire" capture board without the analog-to-digital conversion step. The capture card may also support the delivery of "live" video in addition to "stored" video.

Step 2. Edit/Author: Once the video is converted to digital and is stored on disk it can be edited using a variety of non-linear editing tools. At this stage, as described below, an authoring tool may also be used to integrate the video with other multimedia into a presentation, entertainment, or training format.

Step 3. Encode: After the video is edited and is integrated with other media it may be encoded to the appropriate streaming file format. This generally involves using the encoding software from the video-streaming vendor and specifying the desired output resolution, frame rate, and data rate for the streaming video file. When multiple data rates need to be supported, multiple files may be produced corresponding to each data rate. As an alternative, newer video streaming technologies create one file that has "dynamic bandwidth adjustment" to the needed client data rate.


Step 4. Serve: The video server manages the delivery of video to clients using the appropriate network transport protocols over the network connection. The video server consists of a hardware platform that has been optimally configured for the delivery of real-time video plus video server software that runs under an operating system such as Microsoft Windows NT that acts as a "traffic cop" for the delivery of video streams. Video server software is generally licensed by the "number of streams." If more streams are requested than the server is licensed for, the software rejects the request. Network connections in the Enterprise are generally 10/100BaseT switched Ethernet or Asynchronous Transfer Mode (ATM), while over the public Internet they are IP-based packet-switched networks using dial-up modems, ISDN, or T1 lines.

Step 5. Play: Finally, at the client station the video player receives and buffers the video stream and plays it in the appropriate size window using a VCR-like user interface. The player generally supports such functions as play, pause, stop, rewind, seek, and fast forward. Client players can run stand-alone or can be ActiveX controls or browser plug-ins. They can decode video using software or using hardware add-in decoder boards.

Network Considerations

Not all networks are well suited to the transmission of video. The high bandwidth time-critical nature of video imposes unique demands on network infrastructure and network protocols. Three of the most important characteristics of networks for video transmission are:

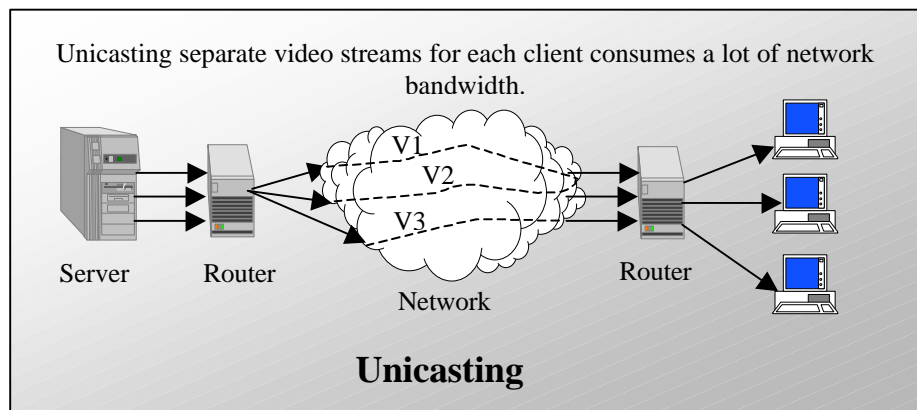
1. High Bandwidth
2. Quality of Service (QoS)
3. Support for Multicasting

High Bandwidth—Digital video ca  Generally, High-Bit-Rate (HBR) video is 1.5Mbps—the rate of MPEG-1—or above, and Low-Bit-Rate (LBR) video is 64Kbps—a single ISDN channel—or less. HBR video needs a high bandwidth network such as a corporate Intranet as a delivery vehicle, while LBR video can go over a network such as the public Internet.

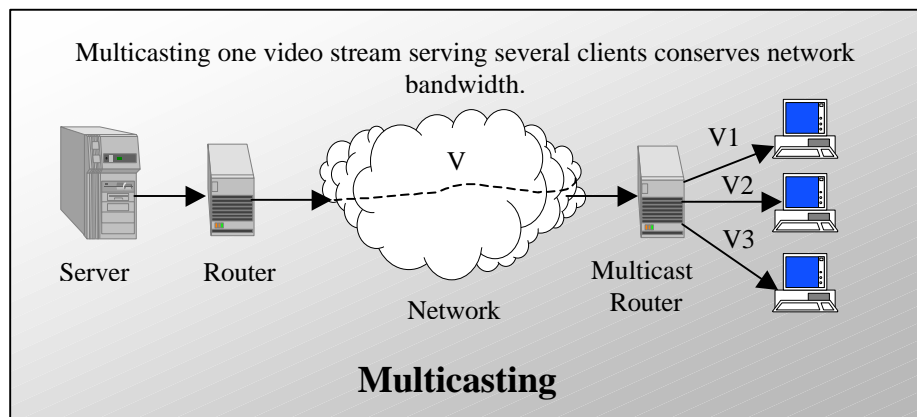
Quality of Service (QoS)—As mentioned earlier, video streaming works best when there is continuous availability of sufficiently high bandwidth in the communications channel. This is especially important for time-critical applications such as video and audio. Variable video/audio packet delays can cause picture jitter and annoying audio starts and stops. To evaluate this for different networks, a new measure of network capability has emerged known as Quality of Service (QoS). Good QoS provides a guaranteed bandwidth at a constant small delay or latency even under congested conditions. Dedicated connections provide the best QoS. Shared media technologies such as Ethernet or the public Internet exhibit variable packet data delays that can play havoc with multimedia data such as video and audio. Specialized protocols and methods have been developed to provide better QoS over these networks.

Support for Multicasting—Another important characteristic for transmission of video over a network is the ability to support multicasting, broadcasting, and unicasting of video streams. Multicasting lies between broadcasting and unicasting.

Unicasting delivers streams 1-to-1 to each client, and is sometimes referred to as "Video-on-Demand" (VoD) since any user can request any stream at any time.



Multicasting, by contrast, delivers streams simultaneously 1-to-many clients where the clients are typically on a subnet of the network. Multicasting is sometimes referred to as "Near-Video-on-Demand" (NVoD) since the subset of users must view the same content at the same time. This is similar to cable TV "pay-for-view" where a subset of users must be authorized to view the program. Applications might include such things as having a subset of users who are all signed up for the same training course. However, there can also be multiple subsets of users viewing different multicasts at the same time..



Broadcasting, on the other hand, is a special case of multicasting that delivers a single stream simultaneously to all network clients. This is the method used for example, for live broadcasts of executive presentations or new product announcements to all corporate employees worldwide. All users on the network can view the broadcast at the appointed time. On the Internet, multicasting is referred to as "IP Multicasting". Since its use has the potential to significantly reduce congestion on the net there is a lot of activity underway to support it. Originally intended for corporate Intranets, the Internet will support it soon as part of the new IPv6 standard. It will require multicast-ready routers, of which there are many already on the market.

LANs/Intranets

Many of today's Enterprise networks were not designed with video transmission in mind. The following is a discussion of various network types and their suitability for video streaming.


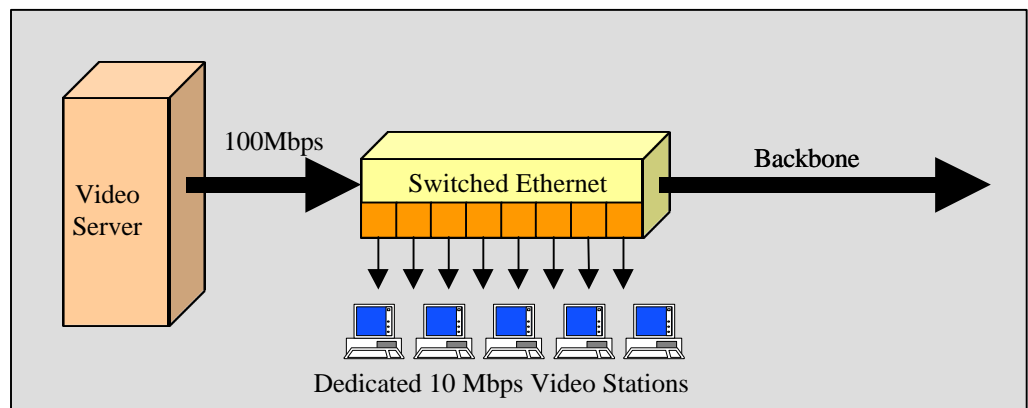
Shared Non-Switched Ethernet

Video streaming can be problematic in a shared network environment such as a 10BaseT shared Ethernet LAN which typically handles intermittent bursty data traffic. This is partly due to two characteristics of Ethernet. First, Ethernet works as a contention-based, or collision-avoidance, system. Each node on the shared section of the network wanting to transmit data listens for collisions with other nodes sending data. If a data collision occurs, all contending nodes back off and re-transmit after a random delay. This enables shared Ethernet to support a large amount of bursty data traffic corresponding to multiple shared applications and devices such as: modem communications, sending and receiving faxes, printing, accessing corporate databases such as phone directories and group calendars, internal and external e-mail, file transfers, and last, but certainly not least, Internet access. However, this coexistence with other network traffic can be disruptive to the smooth playing of real-time video. Non-video-data traffic peaks, for example, could result in lost frames or even stoppage of the video.

In order to handle the demands of video, existing LAN infrastructure may need some modification. Ideally, this can be done without significant cost impact to the customer's investment in existing LAN infrastructure. Some possibilities discussed below include the use of switched Ethernet hubs to provide dedicated bandwidth to a few video stations, Fibre Channel local area digital network, or the use of high-bandwidth ATM for network backbone or long-haul transmission.

Switched Ethernet

A switched Ethernet hub can provide dedicated bandwidth for real-time video to individual desktops or subnets that need it. Generally the switch has a high-speed—e.g., 100Mbps—link to the video server, and multiple lower speed—e.g., 10Mbps—output ports. The switch may also support an optional high-speed link—e.g., Fast Ethernet, ATM, and FDDI—to the corporate backbone.



The Ethernet switching hub only routes video packets that are addressed to the specific client devices on the appropriate output ports. This keeps video traffic flowing smoothly to the desktops that need it, and insulates the corporate backbone network from high bandwidth video.

This establishment of dedicated "subnets" prevents conflicts between video and other data, and also prevents video from hogging shared network bandwidth. Switched Ethernet is often recommended for high-quality video to the desktop as it can often use most of the existing network infrastructure such as Network Interface Cards (NICs), hubs, and cabling.

Fast Ethernet:

Fast Ethernet generally refers to high speed Ethernet such as 100BaseT which supports 100Mbps bandwidth—a tenfold increase over 10Mbps 10BaseT. Fast Ethernet works the same from a protocol and access point of view as standard Ethernet, only faster. The higher bandwidth provides better support for streaming video. Many enterprises install Fast Ethernet as an upgrade or as a network backbone. The next generation of Fast Ethernet—Gigabit Ethernet—is in development.

ATM:

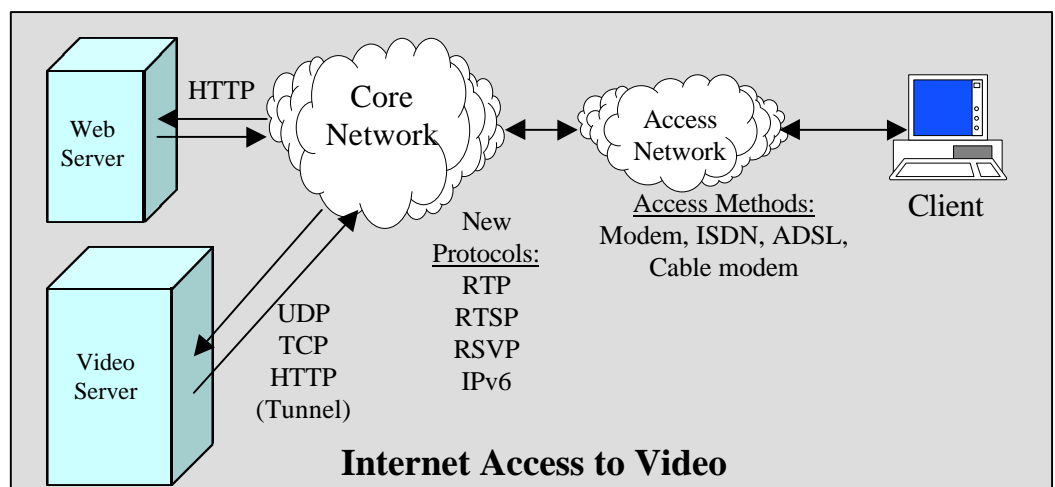
Asynchronous Transfer Mode uses hardware switching to provide not only high speed—25Mbps to 2.5Gbps—but it also a guaranteed Quality of Service (QoS) and multicasting through a circuit-switched connection. It uses fixed 53 byte packets and was designed from the outset for multimedia data transmission such as video, sound, and images. Physical media includes twisted pair, coax, and fiber. Because of its speed and expense, it was originally used for wide-area long-haul networks, but is now beginning to be deployed as a local area backbone.

Public Internet

The public Internet is moving at a breathtaking pace in terms of content, access, and subscribers. There is a strong motivation to be able to deliver video content for a number of potential applications including webcasting, product advertising and catalogs, interactive shopping, and virtual travel.

Access

Access to the public Internet is generally done through an "access network" that connects to the "core network" or backbone of the Internet.



A standard HTTP web server sometimes acts as a "metaserver" that delivers the necessary URLs and information about the video content to the client for access. The video server often uses a different protocol called "Universal DataGram Protocol", or UDP, to provide the client with VCR-like controls and to provide a more "synchronous" data stream with less error-checking overhead than is done in HTTP.

Internet Protocols

As mentioned earlier, the time-critical nature of video makes it difficult to deliver in a normal way over packet-switched networks such as the Internet. For example, standard HTTP web servers using TCP/IP protocol are difficult to use for streaming video. Normal HTTP web pages do not support the 2-way interaction needed to control video streams such as dynamic bandwidth adjustment, rewind, pause, fast-forward, or indexing into the stream. In addition, the TCP/IP protocol adds additional overhead in the form of such things as re-transmissions in the case of missing packets that can be disruptive to a video stream. Most streaming video vendors therefore favor the use of dedicated video servers that use a unique protocol such as UDP (Universal DataGram) for video stream delivery combined with new real-time protocols that have been developed by the Internet Engineering Task Force (IETF) to support streaming multimedia such as video.

IPv6

IP version 6 is a new Internet protocol intended to replace the current IPv4 in use around the world. IPv6 includes extensions for supporting multimedia data streams such as video. It supports multicasting, authentication/encryption, and a larger address space. Since it can't be rolled out everywhere at once, it is intended to be backward compatible with IPv4.

MBONE

The Multicast Backbone is an experimental overlay network on the Internet to gain experience with multicasting that will be supported in IPv6. There are approximately 1700 networks in 20 countries on the MBONE. It is currently being used for multimedia broadcasting such as radio shows and live concerts.

UDP

UDP stands for "Universal DataGram Protocol." Unlike TCP ("Transmission Control Protocol") it does not contain error-checked packets. Hence it avoids the overhead of re-transmission in the case of errors or lost packets. This is OK in streaming video as a few lost packets occasionally will not be noticeable to the eye. The disadvantage of UDP is that it does not penetrate most corporate firewalls. Several vendors have developed work-arounds for this.

RTP

RTP stands for "Real-Time Protocol." It is one of the most commonly used protocols for streaming media on the Internet. It adds a 10-byte header to UDP packets with information such as a time-stamp, sequence number, and type of compression, to enable the synchronized timing, sequencing, and decoding of packets at the destination. This enables multiple media streams such as video, audio, and graphics to be properly synchronized with each other as well as for the packets to be re-assembled in the right order at the receiving destination. Even though packets might arrive skewed by various degrees of network latencies, they can be reassembled and delivered with the timing and sequencing intended. RTP can also be used with other protocols such as TCP or IP Multicast.

RTSP

RTSP stands for "Real-Time Streaming Protocol." It is a higher-level client/server protocol for controlling the efficient delivery and quality of service of a multimedia stream along with controls such as stop, pause, rewind, fast-forward, etc. It also contains provisions for security as well as for usage measurement and rights management so that streaming content providers can control and charge for the usage of their streaming content. RTSP can be used with a variety of IP delivery methods including UDP, TCP, and IP Multicast. It also supports the use of RTP as an underlying protocol.

RSVP

RSVP stands for "Resource Reservation Protocol." Its purpose is to provide consistent quality of service by guaranteeing a maximum allowable transmission delay for streaming packets along with the ability to prioritize specific packet streams. It does this by requesting the dynamic allocation of bandwidth from the network much like PC applications dynamically allocate memory for applications.

Public Broadband Networks

Telecommunications deregulation has enabled local and long-distance carriers to compete with each other as well as with cable TV operators for long-distance services as well as value-added services such as video.


Applications for video over the public broadband networks include:

- ❑ Entertainment—Movies-on-demand, games, and interactive video
- ❑ Education—Courses, libraries, video-clip archives
- ❑ Shopping—retail catalogs, real estate
- ❑ News/Information—News-on-demand, medical care information, sports, weather

Access

Access to the public broadband networks for video is generally through

- ❑ PSTN—using high-speed modems, ISDN, or xDSL technology
- ❑ Cable Modems attached to set-top-boxes
- ❑ Satellite dish receivers with cable or PSTN backchannel



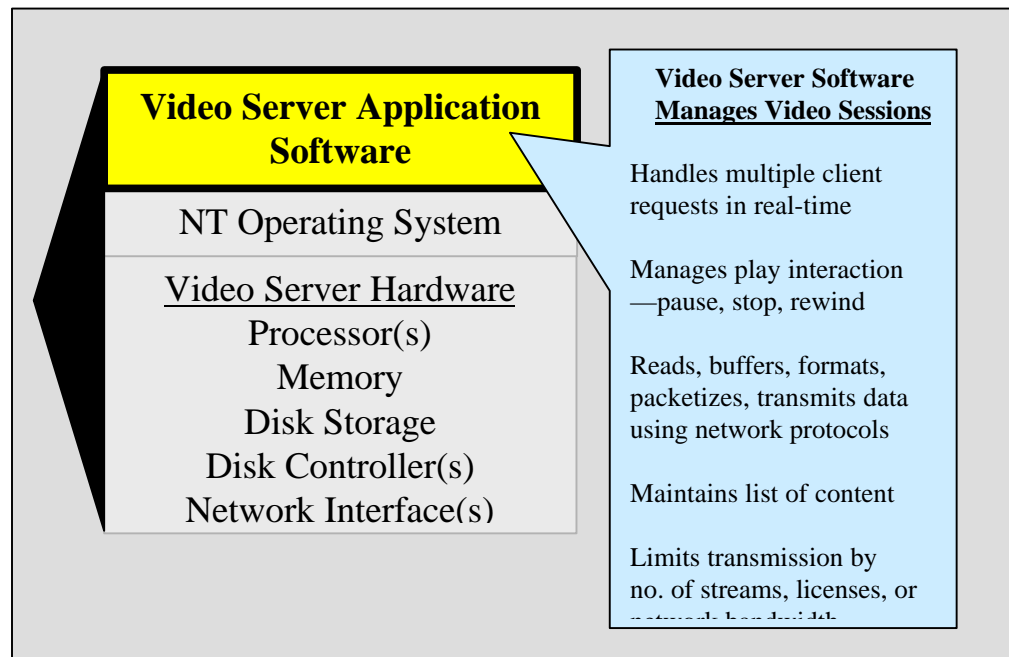
.....

Video Servers

Application Software

The goal of a video server is to match end users to video content, store and retrieve video data as efficiently as possible, transmit it across the network continuously without interruption, and finally, to respond to user interactive VCR-like control. To accomplish this, specialized video server application software runs on a dedicated video server hardware platform that is optimally configured for real-time video delivery. The software running in the video server performs a variety of complex tasks including:

- ❑ Managing video streaming sessions with clients—allocating server resources, limiting transmission by number of streams or network bandwidth consumed, maintaining a list of video content for streaming, refusing sessions that exceed bandwidth or stream licensing limits
- ❑ Managing the playing of individual streams—playing, pausing, stopping, seeking, fast forwarding, rewinding
- ❑ Managing the streaming to multiple clients—sending the same video to several clients simultaneously, or sending different videos to individual clients at the same time
- ❑ Transmitting video over the network—reading, buffering, formatting, packetizing, and sending data over the network using the specified network protocols
- ❑ Storing video data—providing fault-tolerant support
- ❑ Detecting abnormal client termination—e.g., power failure or power off at the client
- ❑ Allowing the loading of new content in real time without interrupting current streams



The client may also perform bandwidth negotiation with the server to determine a preferred bit rate, and therefore which optimized binary file to send to the client or how to adjust a dynamic bandwidth file.

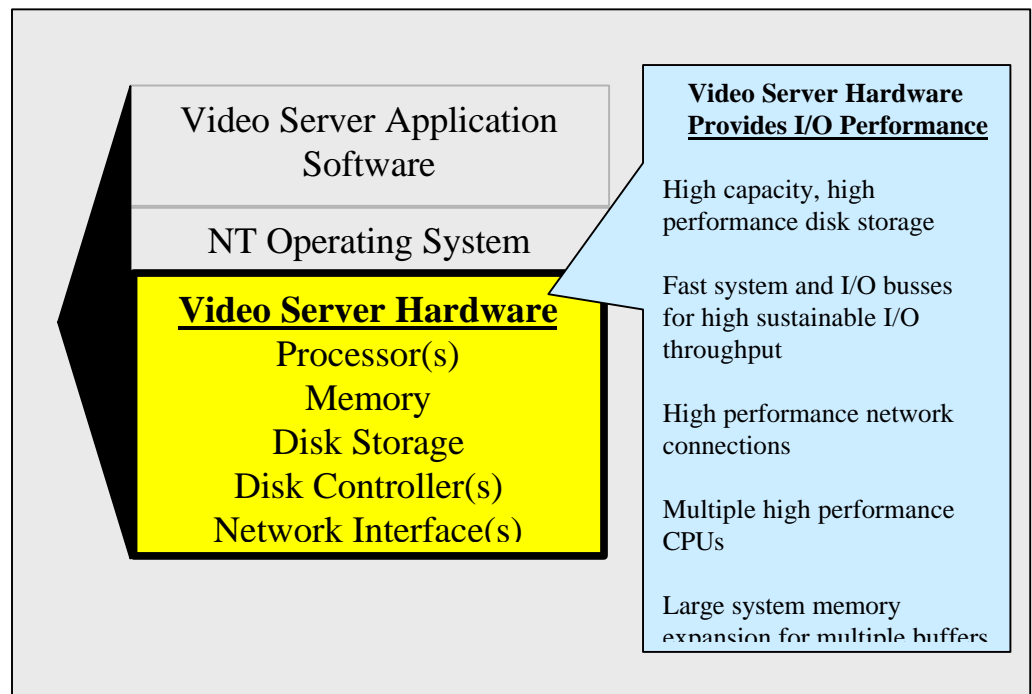
Video server application software that runs under the Microsoft Windows NT Operating System is available from a number of vendors including:

- ❑ Microsoft: NetShow and NetShow Theater
- ❑ Oracle: Oracle Video Server
- ❑ Starlight Networks: StarWorks, StarCast
- ❑ RealNetWorks: RealServer

Video Server Hardware

One of the primary tasks of a video server is to read video data from disk storage, packetize it, and to put it on the network as quickly as possible. This would be a relatively easy task if it were not for the requirement to do this for multiple simultaneous video streams. This key requirement to deliver “real-time” concurrent video streams imposes unique demands on the video server processor, memory, disk storage, and network connections. Video server hardware is typically configured with:

- ❑ High capacity, high performance disk drives for storing and delivering real-time video content
- ❑ High performance network cards such as 10/100BaseT, ATM for high-bandwidth network connections
- ❑ Fast system and I/O busses such as PCI and Wide-Ultra SCSI
- ❑ Tuned device drivers for network cards and disk access
- ❑ Multiple high performance CPUs that can be dedicated to uninterrupted video delivery
- ❑ Large system memory expansion for managing multiple high speed streaming buffer pools



The following is a brief description of some of the most important hardware characteristics to look for in a video server.

High Capacity Disk Storage

As a rule of thumb, MPEG-2 compressed video data takes about 1 GigaByte of storage per hour of video content. A dozen two-hour movies would therefore require about 24GB of disk storage. Therefore, an important characteristic of video servers is their ability to accommodate the highest capacity disk drives such as 9.1 or 18.1 GB. Important parameters to consider include the total number and type of internal drive bays—e.g., the ability to accommodate either 1” or 1.6” disk drives. It is also desirable to be able to expand the storage as the system grows through the use of external storage options using SCSI or Fibre Channel technology. To accommodate all of this disk storage it is also important that there be a sufficient number of PCI slots available for the appropriate disk drive controllers.

High Sustainable Throughput

The real-time nature of video demands that it be delivered to the network as a continuous high bit rate stream. As a result, the most critical resource in a video server is often the I/O bandwidth for providing a sustainable bit rate for multiple concurrent streams. This sustained throughput is usually measured as the average number of bits delivered per unit time. Since video content is generally stored on a rotating disk, this sustainable throughput is greatly influenced by disk rotation speed, disk array architecture, and disk I/O bus throughput. Video content may also be “striped” across multiple disks to improve performance as well as to provide redundancy. It is therefore important that the video server be capable of accommodating high speed drives when needed, as well as wide (32-bit minimum) and fast PCI busses and multiple disk arrays with controllers that support content striping. For example, wide Ultra SCSI-3 disk controllers provide a sustainable throughput of 40 MB/sec—twice that of Fast-Wide SCSI-2. For some applications, the availability of high rotation speed disks—e.g., 10,000 rpm—is also desirable for maximum performance.

High Performance Network

High quality broadcast video can be delivered at bit rates ranging from 3 Mbps to 30 Mbps. This generally requires a high bandwidth network such as 100 Mbps Ethernet connection or a high-speed ATM connection. It is important that the appropriate network interface cards (NICs) be available with highly tuned and thoroughly tested drivers to deliver continuous high bit rate packetized video streams over the network.

Multiple CPUs

Video servers tend to be more I/O intensive than compute-intensive. While they still need to have a relatively fast processor, they don’t necessarily need the fastest CPU available. However, it is sometimes desirable, depending on the capabilities of the video server software, to provide dedicated processors to the “video pump” function. This can prevent other operating system tasks from stealing CPU time from the all-important real-time delivery of video. It can therefore be desirable to be able to expand the server to multiple processors as the system grows.

Expandable System Memory

For smooth delivery, the video server software may maintain multiple buffers in memory for each active video stream to keep data flowing smoothly over the network while the next block of video data is being retrieved from disk. It is therefore desirable to be able to expand the system memory appropriately as the system requirements grow to support more concurrent video streams.

High Availability

It may not be unusual for video servers to be on call to serve up video content at any time day or night to anywhere on the globe. This need for higher than normal system availability may impose

a requirement for optional fault-tolerance. This may come in the form of redundant hardware such as power supplies, fans, or NICs, or in the form of redundant multiple server architectures. To minimize downtime “hot-pluggable” components such as disk drives, power supplies, and fans are also desirable.

Rack-mount for Easy Access

Large video server installations may require multiple servers along with multiple external storage options and networking equipment such as switching hubs. It is desirable to be able to rack-mount this equipment to save space and to provide improved access for maintainability. To accommodate system growth, it should be possible to start out with a small tower configuration and grow to a rack-mount configuration using tower-to-rack conversion features.

Attractive Cost per Stream

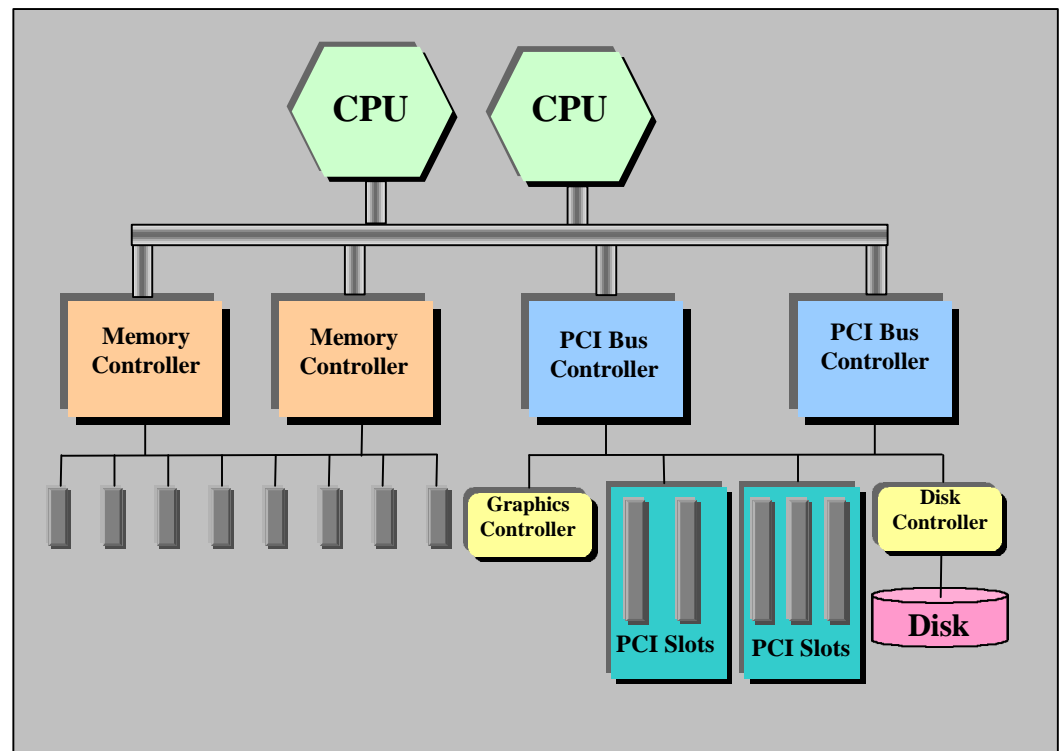
Since video servers are typically sized by the number of concurrent streams they can deliver the “cost per stream” is often an important measure of video server efficiency. This parameter can be a misleading comparison, however, if it fails to take into account the above important hardware characteristics, which also have a cost impact. Nevertheless, the best cost per stream is more likely to come from the ranks of high volume industry-standard architectures.

Example: Compaq ProLiant

The Compaq ProLiant family of servers includes a number of models that make excellent video server platforms. A good example is the Compaq ProLiant 5500. This server features:



- ❑ **Multiple high performance processors.** The Compaq ProLiant 5500 can accommodate up to 4 Intel Pentium Pro 200MHz processors.
- ❑ **Highly Parallel System Architecture for High I/O Performance.** The Highly Parallel System Architecture featured in the Compaq ProLiant 5500 more than doubles the I/O bandwidth needed for demanding video streaming applications.




- ❑ The Highly Parallel System Architecture features:
 - **Dual memory controllers.** Dual memory controllers enable the processing of memory requests in parallel with a memory bandwidth of 1.07GB/sec—up to four times the bandwidth of other systems without this feature.
 - **Dual peer PCI busses.** This enables peripherals on either bus to have independent access to the system's CPUs and memory. The video server I/O load can be balanced by putting high performance peripherals such as disk controllers and graphics controllers on separate busses. I/O bandwidth of up to 267MB/sec—twice that of a single PCI bus—can be achieved. This also means that a larger number of PCI devices can be supported compared to single bus systems.
- ❑ **High performance, high capacity disk storage.** In addition, the ProLiant 5500 supports integrated dual channel Wide-Ultra SCSI-3 disk controllers for maximum sustained video I/O throughput. Internal drive bays accommodate up to 8 hot-pluggable 1" drives or up to 6 hot-pluggable 1.6" with capacities of 9.1GB or 18.1GB per drive. Additional storage is available through the use of SCSI or Fibre Channel external disk storage options.
- ❑ **Large system memory expansion.** The system memory of the ProLiant 5500 can be expanded from 128 MB up to 3 GB for handling high workloads of multiple concurrent video streams.
- ❑ **High availability options.** To maximize availability with minimum downtime, the ProLiant 5500 supports hot-pluggable redundant components such as power supplies, disk drives, fans, and network interface cards.
- ❑ **Large number of expansion slots.** There are 5 PCI and 3 shared EISA expansion slots in the ProLiant 5500 for accommodating such things as additional disk controllers and Ethernet or ATM network interface cards that may be needed in video server applications.
- ❑ **Rack-mount options and conversions.** The Compaq ProLiant 5500 is available in tower or rack-mount configurations to accommodate small or large video server installations. The tower unit features a tower-to-rack conversion capability.
- ❑ **Attractive cost per stream.** By using standard low cost, high-volume components, the Compaq ProLiant 5500 minimizes the cost per stream. From the Microsoft Windows NT Operating System to the industry-standard Intel Pentium processor all the way to the use of low cost high-density Dual In-line Memory Modules (DIMMs) the Compaq ProLiant 5500 provides breakthrough cost per stream for video streaming applications.

Other members of the Compaq ProLiant family that also feature the Highly Parallel System Architecture and other key features needed in video server configurations include the ProLiant 3000 and the ProLiant 1600.

.....
Acronyms

- ADSL—Asymmetrical Digital Subscriber Loop
- ASF—Advanced Streaming Format
- ATM—Asynchronous Transfer Mode
- AVI—Audio-Video Interleave
- CD-ROM—Compact Disk-Read-Only Memory
- CIF—Common Intermediate Format
- CODEC—enCODer/DECoder
- DVD-ROM—Digital Video Disk-Read-Only Memory
- FDDI—Fiber Distributed Data Loop
- GOP—Group of Pictures (in MPEG standard)
- HBR—High Bit Rate
- HDTV—High-Definition Television
- HTML—HyperText Markup Language
- HTTP—HyperText Transfer Protocol
- IETF—Internet Engineering Task Force
- IPv6—Internet Protocol Version 6
- ISO—International Standards Organization
- ISDN—Integrated Services Digital Network
- ITU—International Telecommunications Union
- JPEG—Joint Photographic Experts Group
- LBR—Low Bit Rate
- MBONE—Multicasting BackBONE
- Mbps—Megabits per second
- MJPEG—Motion JPEG
- MPEG—Motion Picture Experts Group
- NIC—Network Interface Card
- NTSC—National Television Standards Committee (U.S. TV standard)
- NVOD—Near Video-on-Demand
- PAL—Phase Alternation Line (European TV standard)
- PCI—Peripheral Component Interconnect
- PSTN—Public Switched Telephone Network
- QoS—Quality of Service
- RGB—Red-Green-Blue
- RSVP—Resource Reservation Protocol
- RTP—Real Time Protocol
- RTSP—Real Time Streaming Protocol
- SCSI—Small Computer System Interface
- SECAM—Sequentiel Couleur Avec Memoire (French TV standard)
- SIF—Standard Interchange Format
- SVHS—Super VHS
- TCP/IP—Transmission Control Protocol/Internet Protocol
- UDP—Universal Datagram Protocol
- VCR—Video Cassette Recorder
- VOD—Video-on-Demand
- XDSL—Digital Subscriber Loop
- YCrCb—Luminance (Y), Color Differences Red – Y (Cr), Blue – Y (Cb)
- YUV—Luminance (Y), Hue (U), Saturation (V)
-
-
-



Resources

For more information on Compaq Computer Corporation, products, and solutions, visit <http://www.compaq.com/>

For information on video server solutions from Compaq visit <http://www.compaq.com/solutions/enterprise/index.html>

For information on standards and specifications affecting video streaming technology visit one of the following:

<http://www.ipmulticast.com>
<http://drogo.cselt.stet.it/mpeg/>
<http://www.iec.org/>

Multicasting information
MPEG home page
Internet standards documents

For information on specific video server application software visit one of the following web sites:

<http://www.microsoft.com/theater/>
<http://www.starlight.com/starlight/start.htm>
<http://www.oracle.com/products/asd/video/video.html>
<http://www.realnetworks.com/>

Microsoft NetShow Theater Server
Starlight Networks video server software products
Oracle Video Server
RealVideo Server

For further reading on video technology and video streaming consider the following:

Books

- Video and Media Servers, Technology and Applications, Focal Press, Karl Paulsen, 1998
- Publishing Digital Video, Academic Press, Jan Ozer, 1997
- A Technical Introduction to Digital Video, Charles A. Poynton, John Wiley & Sons, 1996
- Digital Video: An Introduction to MPEG-2, Haskell et. al., Chapman & Hall, 1997
- Video Dialtone Technology, Daniel Minoli, McGraw-Hill, 1995
- The Digital VideoMaker's Guide, Kathryn Shaw Whitver, MWP, 1995
- Video Conferencing and Interactive Multimedia, Toby Trowt-Bayard, Flatiron Publishing, 1997

Articles

- *Streaming Video: The Time is Now*, Jan Ozer, Emedia, April 1997
- *IP Multicasting Expands the Web Pipeline*, David Kosiur, ZD Internet Magazine, August 1997
- *Video on the Web: Server vs. HTTP Video Streaming*, Nels Johnson, DV Magazine, October 1997
- *The Complete Video Desktop*, Jan Ozer, PC Magazine, October 7, 1997
- *Getting Video on the Web*, Nels Johnson, DV Magazine, August 1997
- *The Real Thing*, Nels Johnson, DV Magazine, June 1997
- *Live Video on the Web*, Louis M. Brill, DV Magazine, August 1997
- *Watching the Web Stream By*, William Wong, ZD Internet Magazine, September 1997
- *Digital Television: Making it Work*, Bhavesh Bhatt et. al., IEEE Spectrum, October 1997