DIGITAL VIDEO
FOR THE
NEXT MILLENNIUM

INTRODUCTION

At the end of 1998, networked digital video was very much a work in progress. Digital video files could be created, shared and stored but not in the robust, transparent manner that computer users expect and receive for other applications. Digital video client/server systems supported modest implementations but did not scale to adequately support high-bandwidth traffic or shared services among multiple locations. A critical issue for digital video client/server systems was their inherently proprietary nature. Moving assets from one system to another could require re-encoding of video assets, or at the very least, re-authoring of Web pages providing access to the video assets. File format support required specific encoder systems and decoders at the client. Asset management and file indexing capabilities were generally weak. Digital video vendors are very committed to product development, with the result that innovative services are announced almost daily. The state of the art is fluid and exciting.

Five institutions--University of North Carolina, Chapel Hill, University of Tennessee, Knoxville, Georgia Institute of Technology, NYSERNet, and North Carolina State University, established the Video Development Initiative (ViDe) in 1998 to identify and prioritize digital video development issues. This ambitious goal involved issuing an RFI to vendors to identify and evaluate the existing state of the art and then selecting partners among responding vendors to develop highly scalable, functional, interoperable and robust access to video resources, across platforms, client/server systems, institutions and countries.

The Southeastern University Research Association (SURA) awarded ViDe a grant for Phase I of this multi-year initiative. A key Phase I objective was providing information and support for colleagues implementing digital video projects. Several Phase I deliverables addressed the need for digital video information in the academic community. The *Video Conferencing Cookbook* (http://sunsite.utk.edu/video_cookbook/) was released in February, 1999. A SURA and ViDe sponsored workshop was held in Atlanta, Georgia in March, 1999 (http://vide.utk.edu/conference/agenda.html) This white paper, *Digital Video for the Next Millennium* (http://sunsite.utk.edu/video/) provides an overview of digital video on demand--the underlying technology, the client/server capabilities currently available and development areas for the near future. In the coming years, ViDe will partner with selected digital video vendors to develop critical capabilities for highly functional and available digital video in the academic community.

SECTION ONE: THE DIGITAL VIDEO PROCESS

The simplest definition of digital video is the representation or encoding of an analog video signal in digital bits for storage, transmission and display. If you have access to the World Wide Web, chances are you have viewed digital video files. Examples include CNN news clips, movie trailers and, of course, the popular dancing baby! Digital video files pervade the analog television and 70 mm film world. Most special effects, such as Godzilla trampling a building or a polar bear sipping a soft drink in an advertisement, are created by editing digital video files. If you have rented movies on demand in your hotel room, played a DVD, a video game, or used a Direct TV satellite dish; you have experienced digital video through your television set.

Digital video is a growing presence in the academic arena, from digitized course lectures to archival footage housed in the campus library. Video conferencing--for collaboration, Internet-based communication and teaching--is an important digital video service. Video conferencing is addressed by ViDe in the Video conferencing cookbook. Digital video on demand, another key service, is defined for this white paper as the creation, storage, transmission and display of archived digital video files in an IP network environment. Digital video may be streamed to a computer, so that playback begins before the entire video file is received, or non-streamed, requiring that the entire file be downloaded be playing. Streaming videos may be served as multicast or unicast streams. Video on Demand generally refers to unicast, where a single video file is requested by a user and streamed to the user's computer for playback. Multicast--the transmission of a single digital video file to multiple users in a scheduled environment, is included in this digital video on demand white paper for convenience, since most vendors providing on-demand video files to a single user (unicast) also provide products for multicast of both stored files and live broadcasts.

There are three basic components in the digital video process.

Step One: Encoding or Digitizing the Video

The analog signal, which can be a direct broadcast or a videotape, such as a VHS cassette, must be digitized. An encoder card accepts an analog signal through a cable into an interface card and feeds the signal into the encoding hardware and software to encode the video into digital form.

Encoding is a simple concept: the video analog signal is encoded, or represented, in digital bits that can be read and operated upon by a computer processor. All digital files--whether a textual document, an image, a program, or a video--are representations of information in bits.

One meaningful element in graphical digital media is the pixel, or picture element, which is a two-dimensional base unit of programmable color, also represented in bits. A pixel is a logical element, so that its representation can vary based on factors such as bit depth and screen resolution. The color expressed by a pixel is a blend of some component of the red, green or blue color spectrum. The human eye has photoreceptor cone cells that respond to color in the three spectra, so three mathematical representations of color--red, green, and blue--are all that are needed for digital color reproduction. Like all other digital information, colors are represented by bits. More bits (8-bit, 16-bit, 24-bit, etc.) allow more precise definition of the exact blend or hue from the red, green, blue color spectrum. Analog colors are translated to pixels in the RGB digital color space. Digital video uses a non-linear variation of RGB called YCbCr, where Y represents luminance, or brightness, and CbCr represents chrominance (chromaticity), or "pure" color, in the absence of brightness.

The number of pixels displayed on your computer screen, along the horizontal axis and the vertical axis, is defined as the spatial resolution. Broadcast-quality digital video (CCIR 601) is commonly displayed with 720 x 480 resolution.

Video is more than color, however. Video requires multiple frames showing redundant information and fractional changes to create the illusion of motion across time and space. At some point in your childhood, you probably duplicated the illusion of motion by drawing stick figures on cards and flipping them to create your own "cartoon." The more redundancy between frames, the smaller the change from one frame to the next, the smoother and more continuous the illusion of motion on your television or movie screen. Video encoding algorithms take advantage of this redundancy to compress video, encoding only the difference between frames. This process is known as temporal compression. The decoder at the client end stores the information that does not change from frame to frame in the buffer to refresh the displayed frame as needed.

To convert analog video to digital, each frame must be digitized using encoding hardware and software. Encoding systems can be as inexpensive as a $300 card fitting into an available slot on a multipurpose microcomputer to a $10,000+ stand-alone system, which requires a dedicated microcomputer. A good tape deck and analog monitor are also usually required for the encoding process, depending on the requirements of the selected video card/encoding system.

A video-encoding card accepts analog inputs from a VCR or a video camera and converts the analog format into a digital video file. Encoding hardware and software vary greatly in cost and therefore support a wide range of functionalities, including, as the cost increases, higher quality output, separate input for video and audio, faster encoding, multiple file and batch file processing, analog output (e.g. digital video back to analog videotape), uncompressed conversion and, at the present time, a range of encoding formats, including M-JPEG, MPEG-1, editable MPEG, MPEG-2, Video for Windows/ActiveMovie, and QuickTime.

Video cards are available for proprietary formats such as Intel's Indeo®. Video cards create digital files that can be opened by editing software, such as Adobe Premiere®. Video editing packages allow you to make changes to digital video, such as adding credits or special effects, cutting or adding frames, merging digital video clips, and outputting the created movie to a range of digital file formats, for playback in a variety of ways, through the use of incorporated software or plug-ins.

When digitizing video, each frame must be converted to digital form. In addition, the audio track accompanying the digital video must be converted and synchronized to the video for playback. A straight digitization with no compression requires more bandwidth and processing power than desktop computers can handle. Currently, compression must be employed to convert analog audio and video so that it is usable at the desktop. Compression reduces redundant information so that meaning is not lost but file sizes are reduced to manageable form. When you "rent" a movie in a hotel room, you are seeing a slightly-compressed (MPEG-2) version of a movie you might have seen in uncompressed analog form a few months earlier at a movie theater.

Uncompressed digital video can be created on high-end platforms and broadcast in real time or stored for later use. The Society of Motion Picture and Television Engineers (SMPTE) and the Institute of Electrical and Electronics Engineers (IEEE) develop and manage standards for uncompressed digital video. These standards include CCIR-601 for television broadcast digital video, in the resolution of PAL, NTSC, and SECAM; SMPTE 259M for the transport of CCIR-601; and SMPTE 292M for the transport of high definition television (HDTV).

The Moving Picture Experts Groups, known collectively as MPEG, are responsible for developing and maintaining digital video and audio encoding standards to address a wide range of commercial and educational needs. MPEG employs established procedures for the development, adoption, testing and review of digital multimedia standards. Standards are published and made freely available to commercial developers, although reasonable costs for some technologies may apply. MPEG standards are international standards that insure video encoding systems will create standardized files that can be opened and played at any desktop with a standards-compliant decoder. MPEG encoding standards are discussed in more detail in Section 2.

In the past few years, digital video cameras have become available, in commercial and consumer-quality models. A high performance serial bus, IEEE P1394, popularly known as FireWire, was developed by Apple Computer but now supported by many vendors to support data transfer rates of 100, 200 or 400 Mbps.

These high transfer rates mean that digital video can be transported directly from the digital source (camera, DVD, etc.) into the microcomputer with no processing delays. FireWire streams video data off a hard drive in real time without computer assistance. FireWire supports up to 63 devices on a single bus, which can be connected in a star, tree, or daisy chain pattern, and allows 1,023 buses to be bridged together. Addressing is dynamic and allows devices to be connected without rebooting the computer.

FireWire transfer speeds, currently at 100-400 Mbps, will increase to 800 Mbps/multi-Gbps in the next release--1394B. The high transport speeds can result in latency problems, requiring significant buffering capacity, as can a heavily-loaded PCI bus, but these problems will abate as FireWire integration becomes the norm, and microcomputers are designed for FireWire integration.

One issue with digital video camera content creation is that the resulting digital video files are generally in formats not currently supported by digital video client/server systems. When investigating digital video cameras, insure that a method exists to output the file to a standard format (AVI-to-MPEG or directly to MPEG)--whether through an editing program such as Adobe Premiere or a transcoding system such as Heuris.

Step Two: Sending Digital Video to the Desktop

Once a video is created, it is stored and then transported to the desktop for playback. Digital video created on a computer can be stored on the computer, opened and played back, just as a document is opened in a word processing program for reading, editing and printing.

A server must generally be employed to store and share a video over a network--whether a campus or building LAN or the Internet. Digital service includes real time broadcast, non-streamed downloading or streaming to the desktop. Video service may be multicast (“one to many”) where one video stream is served to many viewing clients or unicast (“one to one”) where one video stream is served to one viewing client.

Real time broadcasting converts analog video to digital on the fly. Analog video is received by the video server directly from a broadcast feed or a video camera, encoded in real time, and then served as a multicast video stream to many clients. Real time broadcasting also includes the real time delivery of files already in a digital format, such as a digital camera or satellite transmission.

Video files are meaningful only when forward progression, providing continuity of information, is maintained. A cartoon coyote cannot be running off the cliff in one frame and standing on the edge of the cliff, looking down, in the next frame, if the video file is to make sense to the viewer. The coyote also cannot be running swiftly to the cliff in one frame and moving slowly and jerkily in the next. Video data must be played in the correct order, with little or no packet loss, and with smooth, continuous timing or else essential information will be missing. To insure that video files are usable at the desktop, the frames must be received in order and timed for playback. Digital video can be received at the desktop for playback in two ways: non-streaming or streaming video.

Non-streaming video requires that an entire video file be downloaded and lacks the timing functionality for smooth packet streaming. A video server is not required to store and serve non-streaming digital video.

Any server can store and serve non-streaming video, or the non-streamed file may be stored on the microcomputer hard drive or a CD-ROM and played back.

Digital video functionality, for opening the non-streaming file and playing it back on the client machine, is provided by the client software. Downloaded video is an option when the latency (elapsed time) required for the download process, which can range from several minutes to more than an hour, is not an issue.

Non-streaming video is also employed when a video server is not available to provide streaming. Non-streaming video files may also be provided when the maximum number of concurrent video streams supported by a video server has been exceeded. Most video server vendors do not support download of non-streaming videos. If an institution wants to offer nonstreamed video files for download, to insure high availability for the files, an FTP server or other download site must be separately provided.

Streaming video begins playback on the client as soon as enough of the video has loaded to begin and sustain playback at a continuous rate. Cache is established from random access memory (RAM) on the client desktop and is used to receive the file, insure that frames are in the correct order, establish timing, refresh compressed frames and check for dropped packets. The video file continues to download into the client cache even as the beginning of the video is being viewed. Video streaming relies on technology at the video server and at the client, such as caching and control bits, to receive and assemble a video in which all data bits play smoothly, in progressive frame order.

Video streamed via the Web must be transported within the IP architecture. Streamed video has low tolerance for the enforced reliability of TCP, which would keep an application waiting for the retransmission of dropped packets. UDP (User Datagram Protocol) is frequently used in place of TCP as a transport protocol for real time applications, such as digital video.

UDP uses the Internet Protocol (IP) to transport a data unit ("datagram"). UDP supports digital video because it does not divide the data stream into packets for reassembly at the client end. However, UDP also does not order the datagrams into the correct sequence. Applications using UDP must insure, at the receiving end, that the complete message has arrived, in the correct sequence order.

RTP (Real-time Transport Protocol) is a UDP protocol that provides payload type identification, sequence numbering and time stamping. RTP allows for packets to be transported out of order and reassembled in correct order at the receiving end. Digital video has low tolerance for disordered packets and dropped frames. It is used on the MBONE, for interactive audio and video, particularly conferencing sessions. RTP is used with a companion protocol, RTCP (Real-time Control Protocol), which provides periodic control packets to an application to monitor the quality of the data distribution.

RTSP (Real-time Streaming Protocol) is an application-level rather than a simple protocol, since it works with many transport protocols--TCP, UDP, RTP, and IP Multicast. RTSP was designed to support streaming multimedia in unicast and multicast applications. It provides increased functionality at the client end for playback, seeking, etc. and has been described as a "video remote control" for the computer. Among other features, RTSP allows for interoperability between server and client implementations from different vendors. RTSP can be used with RSVP to establish and manage reserved-bandwidth streaming sessions. Progressive Networks' Real Player G2 is an example of an RTSP client.

RSVP (Resource Reservation Protocol) provides Quality of Service (QoS) by allowing an application invoking RSVP to reserve end-to-end bandwidth, memory and CPU resources sufficient for the demands of the application. RSVP requires that all network components work together to provide guaranteed resources for the application, so all components--hosts, routers, hubs, etc.--must support RSVP. Although RSVP is a fairly mature standard, it is not heavily implemented, due at least in part to the requirement that all network components support the protocol.

IP Multicast supports one-to-many service for a data stream. All routers in the network infrastructure must be IP multicast-enabled. Some multicast applications, such as Progressive Networks' RealSystem G2, are able to bridge non-multicast-enabled network segments. A process asks its host for permission to join or leave a group. IP multicast-enabled routers query their groups to identify the processes currently belonging to each group. On the Internet, IP Multicast is implemented on the MBONE and, increasingly, by service providers who will multicast your video to your authenticated users for a fee. Multicast is supported natively on advanced networks, such as vBNS and Abilene.

Quality of Service (QoS) provides a mechanism whereby a client can request priority access, sufficient bandwidth and other network service characteristics to guarantee acceptable application performance. QoS sounds deceptively simple but is difficult to implement since the QoS protocol (such as RSVP) must be enabled on all network devices to insure that the bandwidth allocation is supported across the network. QoS becomes increasingly important as more clients request unicast video applications or participate in a multicast transmission. Most networks currently utilize admission control (users beyond a prescribed stream limit are denied service until a stream is freed) or "best effort," where all users share bandwidth equally and suffer equally, as bandwidth utilization approaches capacity.

Step three: Playback at the Desktop

When a streamed video file is received at the desktop, the file type must be recognized, through information provided in the header, and then opened. The file must be cached until sufficient data is received to allow for smooth, continuous playback. Playback includes controls such as forward, reverse, stop and play, as well as freeze-frame, content bookmarking, audio volume control and sizing of the viewing window. Client software, frequently a web-browser plug-in or browser helper-application, provides this functionality. If you visit a web site with audio or video, you are usually alerted to the plug-in viewer needed and provided a link for free download of the plug-in. RealVideo, Windows ActiveMovie (AKA NetShow), and QuickTime are common file formats with freely-available viewers.

In the past, different client players were required for different file types. Currently, many standard client players have extended their client capabilities to recognize multiple standard file types, such as MPEG-1 (.mpg), ActiveMovie/Video for Windows (.avi), QuickTime (.mov) and RealVideo (.rm and .ram). Progressive Networks' G2 Real Player, with plug-in extensions, and Microsoft's Windows Media Player can open and playback multiple digital video file formats, for example.

Additional functionality provided by the client, in collaboration with the server, can include permission to save a file to hard drive, optimization of playback based on network connection (e.g. 28.8 KB, 56 KB, T1, etc.) and auto-selection of video file format for a video asset transcoded in multiple formats, usually based on the bandwidth capability of the client network connection.

SECTION TWO: VIDEO ENCODING STANDARDS

The critical first step in the digital video process is the encoding of analog video to digital form. As mentioned in the last section, this usually requires compression of some information to create a usable digital file. There are many encoding technologies that create digital files that can be stored, served and played back on PC and Mac desktops. The only international standards--developed and ratified according to established procedures by an authorized maintenance agency--are those which have emerged from committees of the Moving Pictures Experts Groups: MPEG-1, MPEG-2 and MPEG-4.

MPEG-1 (ISO/IEC 11172)

The first digital video and audio encoding standard, MPEG-1, was adopted as an international standard in 1992 to provide digital video at bit rates up to 1.5 Mb/sec.* The impetus for the standard was to provide encoding and playback of VHS-quality digital video for CD-ROM playback. MPEG-1 is a progressive-video-sequence encoding standard. The standard implementation for MPEG-1 (known as “constrained bit stream”) supports 352 pixels x 240 lines/sec at 30 frames/sec and requires 1.5 Mbps bandwidth for transport. MPEG-1 compression relies on the considerable redundancy of information within and between frames to compress a video object without significantly compromising the integrity of the information it contains.

Video contains spatial, spectral and temporal redundancies, which may be compressed without significant sacrifice in meaning. The encoding techniques in MPEG-1 involve compression based on statistical redundancies in temporal and spatial directions. Spatial redundancy is based on the similarity in color values shared by adjacent pixels. A red sweater in a video frame will generally possess a uniform color value, with little or no perceptual variation from one pixel to the next. MPEG-1 employs intraframe spatial compression on redundant color values using DCT (discrete cosine transform).

Spectral redundancy in video is the similarity between color spectra or “brightness.” MPEG-1 operates in the YUB color space. RGB data is converted to YUB. 24-bit RGB is subsampled at 4:2:0 YCrCB, where Y = luminance (brightness) and CrCB = chrominance (color difference). The human eye distinguishes difference in brightness more readily than difference in pure color value.

.

* The standard actually scales higher than 1.5 Mb, but 1.5 Mb is the accepted "sweet spot" for MPEG-1

Temporal redundancy is the sameness in temporal motion between video frames. If frames were not redundant, there would be no perception of smooth, realistic motion in video. MPEG-1 relies on prediction--more precisely, motion-compensated prediction--for temporal compression between frames. MPEG-1 utilizes three frames to create temporal compression—I-Frames, B-frames and P-frames. An I-frame is an intra-coded frame, a single image heading a sequence, with no reference to past or future frames. MPEG-1 compresses only within the frame with no reference to previous or subsequent frames. P-frames are forward-predicted frames, encoded with reference to a past I- or P-frame, with pointers to information in a past frame. B-frames are encoded with reference to a past reference frame, a future reference frame or both. The motion vectors employed may be forward, backward, or both. B-frames are also sometimes known as digital video "spackle."

The MPEG-1 coding standard is a generic standard, intended to be independent of a specific application, serving as a toolbox to be adapted to different applications and their associated hardware and software.

MPEG-2 (ISO/IEC 13818- 2

MPEG-2, published as a standard in 1994, is a high-bandwidth encoding standard, supporting a bandwidth range of approximately 2Mbps to more than 20 Mbps. It was originally designed for coding of television broadcast video with CCIR Rec. 601 resolution at data rates below 10 Mbps, but was expanded to encompass HDTV requirements at app. 12-20 Mbps.

MPEG-2 was designed to encompass, and be backward compatible with, MPEG-1 encoding techniques but was also enhanced to support interlaced video, as provided by television input sources. The MPEG-2 standard was designed for scalability and flexibility, supporting many levels of service depending on the needs of the application. It was expected that an MPEG-3 standard would be developed for HDTV (high definition television), but the MPEG-2 standard scaled to encompass the bandwidth requirements of HDTV.

The most common MPEG-2 compression is main level ("CCIR 601") at 720 pixels x 480 lines, 30 frames/second. The sweet points for MPEG-2 support the bandwidth bit rates of 2-6 Mbps, scaling up to 40 Mbps for very high-level HDTV applications.

The MPEG-2 encoding standard builds on, and is backward compatible with, the statistical redundancy compression of MPEG-1. The most important difference between MPEG-1 and MPEG-2 is the encoding of interlaced frames for broadcast TV. MPEG-1 supports only progressive frame encoding, while MPEG-2 provides both progressive frame and interlaced frame encoding. Video movies, originally in a film format, are a progressive frame format.

Television broadcasts are an interlaced format. A broadcast frame is created with two separate fields, a top and bottom interlaced field, with the first line of the bottom field appearing immediately after the first line of the top field. MPEG-2 splits frames into two fields for interlacing, so that 30 frames/sec becomes 60 fields/sec.

In addition, MPEG-2 includes the ability to multiplex video streams, additional color subsampling, improved compression and error correction and improved audio, including "low sample rate" and multichannel extension for surround sound. The many profiles and levels of service include NTSC (app. 3Mbit/sec), PAL (app. 4 Mbps) and Broadcast HDTV (12-20 Mbps).

MPEG-4 ISO/IEC 14496

MPEG-4, the latest encoding standard from MPEG, was finalized in October 1998 and should be ratified as a standard in the first half of 1999. MPEG-4 arose from a need to have a scalable standard supporting a wide bandwidth range from streaming video at <64 Kbps, suitable for Internet applications, to app. 4 Mbps for higher-bandwidth video needs. MPEG-4 also arose from a desire, as digital encoding matures, to advance beyond simple conversion and compression to object recognition and encoding, as well as the provision of synchronized text and metadata tracks, to create a digital file that carries a meaning greater than the sum of its individual parts.

MPEG-4 supports both progressive and interlaced video encoding. The standard is object-based, coding multiple video object planes into images of arbitrary shape. Successive video object planes (VOPs) belonging to the same object in the same scene are encoded as video objects. MPEG-4 supports both natural ("analog") and synthetic ("computer-generated") data coding. Some VRML technology is incorporated to encode dimensionality.

MPEG-4 compression provides temporal scalability utilizing object recognition, providing higher compression for background objects, such as trees and scenery, and lower compression for foreground objects, such as an actor or speaker—much as the human eye filters information by focusing on the most significant object in view, such as the other party in a conversation. Object encoding provides great potential for object or visual recognition indexing, based on discrete objects within a frame rather than requiring a separate text-based or storyboard indexing database. In addition, MPEG-4 provides a synchronized text tract for courseware development and a synchronized metadata track for indexing and access at the frame level.

Proprietary Formats

MPEG-1 and MPEG-2 were developed for reliable moderate to high-bandwidth transport. Neither standard successfully supports streaming over the Internet, particularly at the common modem speeds of 28.8 to 56 Kbps available for personal Internet use. Encoding formats proprietary either to a microcomputer platform or a specific manufacturer arose to provide streaming digital video to Internet users at low-bandwidth ranges. MPEG-4 is expected, over time, to displace many proprietary formats but at present, proprietary "de facto" standards are well-established on the Internet.

QuickTime began as a Mac-based video encoding, file management and playback system but with version 3.0 became a cross-platform encoding format, supporting digital video on Mac and Windows. QuickTime video files have the file extension .mov. QuickTime is a versatile digital video encoding format, supported by a range of commercial and shareware software products, including encoding, editing, and client plug-ins. QuickTime 4.0, released in April, 1999, supports timecode tracks and Web transport and streaming protocols, including HTTP, RTP and RTSP. QuickTime 4.0 provides built-in support for digital video, including MiniDV, DVCPro, and DVCam camcorder. The QuickTime digital video file format was selected as the basis for MPEG-4. Although the multi-track, object-based MPEG-4 goes beyond QuickTime functionality, the QuickTime wrapper will be supported by MPEG-4-based service and streaming, making this a safe interim choice for low-bandwidth videos over Internet.

Microsoft ActiveMovie (AKA Video for Windows), provides digital video technology, including encoding, file naming and playback, on a Windows platform. The file format is identified by the .avi (audio & video interleave) extension. Video and multimedia creation and editing packages for use on Windows platforms, such as Adobe Premiere® must support the .avi file format

Microsoft ASF (Advanced Streaming Format) is an open streaming format developed collaboratively by Microsoft, Progressive Networks, Inc., Intel Corp., AdobeSystems Inc. and Vivo Software Inc., as well as from the feedback and suggestions of other companies. It is currently available in beta in version 1.0 for use with Microsoft's NetShow streaming media client/server software and in version 2.0 as a preliminary developers' toolkit. The ASF format is intended for streaming synchronized audio, video and multimedia for use over the Internet.

Progressive Networks' RealVideo (file format extension .rm or .ram) is a robust low-bandwidth format intended for Internet streaming. RealVideo is supported by a range of commercial and shareware products including encoding software, multimedia authoring tools, server software and client plug-ins. The G2 Real Media client supports SMIL synchronized text files, audio only, and plug-in extensions for MPEG. RealVideo is in widespread use and is supported by a complete client/server suite, including publishing, synchronized multimedia, and streaming server software, available in shareware and inexpensive commercial versions.

Other encoding formats proprietary to different manufacturers include Intel’s Indeo®
format and **Cinepak**, first developed by SuperMac Technologies and now owned by Radius.

M-JPEG is a quasi-standard provided by many video encoding cards. M-JPEG consists of sequential JPEG-encoded frames. JPEG stands for Joint Photographic Experts' Group and is the popular name for the still image encoding standard JFIF (JPEG File Interchange Format). JPEG is an intraframe compression standard intended for still images only. Video encoding cards provide M-JPEG compression so that the resulting digital video file may be edited. MPEG-1 and MPEG-2 files include I-frames, which are compressed intraframe, but also P-frames and B-frames, which do not include essential information, such as color or movement, but instead reference that information in a forward or backward frame. Since frame types are not eye-readable, video editing can result in the removal of critical reference frames. Some video cards now provide editable MPEG-1, eliminating the need for M-JPEG encoding. Editable MPEG-2 was introduced by Hewlett-Packard in 1997 for its HP MediaStream broadcast server.

SECTION THREE: SELECTING A DIGITAL VIDEO CLIENT/SERVER SYSTEM

Digital audio and video exist today in a state of paradox: demand for digital files in audiovisual formats is great; mature standards supporting digital video and audio on optical media (CD, CD-ROM and DVD) and broadcast systems (satellite, TV, HDTV) are widely available. Systems to support digital video over IP, however, while offering exciting functionalities and tremendous potential, are still emerging rather than proven, particularly in critical areas of interoperability and adherence to open standards. It is possible to select the right system for your needs, among many good, fairly solid offerings, but more care must be taken than in the selection of client/server systems to support electronic text. In early 1999, selecting the right digital video client/server system to grow with your developing needs is still a risky venture, requiring careful planning and needs assessment.

Hardware and software must be selected for two separate but convergent processes to support digital video: encoding (file creation) and service (file storage, transmission and display). Generally, these two processes are completely separate purchases, yet they must seamlessly interoperate. A major problem with digital video, even in 1999, six years after the adoption of the MPEG1 standard, is the fact that most digital servers remain sensitive to the digital encoder card used to create MPEG1 and MPEG2 files, as well as to the decoder cards used to read MPEG files, particularly MPEG2 files. The situation improves each year but remains a frustrating and expensive problem to solve, depending on the systems selected. Lack of interoperability can involve custom API programming, particularly at the client end, serious performance issues at the client, and, in some cases, re-encoding of the video assets.

How can this problem be addressed? If at all possible, the encoding system and the client/server system should be selected in tandem or at least with due consideration to the interoperability of both processes. Don't encode too many files before knowing that the encoding system selected can be supported by the chosen client/server system. Otherwise you may be limiting your options to a client/server system based on the encoding system rather than on user need. It is tempting to say that the more expensive system should be chosen first. But which system is more expensive? In actual dollars, a client/server system will generally range from $15,000 to $100,000 while the encoding system ranges from $300 to $10,000. However, video encoding is a time-consuming, labor-intensive process. The cost in staff hours is much greater for encoding than for client/server set-up and support. Also, like all magnetic media, videos have a fairly short shelf-life, which decreases, with each use. Rare, irreplaceable videos should be encoded only once, and then digital use copies should be struck from the digital master.

It is best to give precedence in selection to the client/server system. The client/server system provides the functionality to support current user needs and to respond to changing need. Therefore it has the greatest impact on user satisfaction. Any encoder system supporting MPEG standards at an acceptable level (composite bit stream for MPEG-1, main level for MEPG-2) is going to produce acceptable files for the user, provided the operator creating the files is competent and well-trained. The human eye cannot distinguish small differences in quality among MPEG files. Problems with the client/server system, such as denial of service, jitter and poor audio/video synchronization resulting from streaming problems, incompatible files, client interoperability issues, or network overload will be very noticeable to the human senses.

When selecting a client/server system, load a copy of the client on a range of local computers and test MPEG1 files created from different encoding systems. These files may be created in-house, "borrowed" from other institutions, or found on the web. For example, the Library of Congress now offers MPEG videos for view. (add URL). Test different frame resolutions, bandwidth encoding speeds and, preferably, 30 frames/second. Test files with audio (talking heads, music) and video, high action, etc. Compare playback quality among vendor clients and also compare playback quality against inexpensive or shareware players, such as Window Media Player, VMPEG-1.7 from the MPEG Software Simulation Group, or Xing MPEGPlayer.

It is possible, and sometimes unavoidable, to select the encoding system first. Most client/server vendors test encoder cards for interoperability and either publish acceptable encoder systems on their web pages or provide that information to prospective buyers on request. A fairly safe strategy is to purchase an encoding system supported by the largest number of client/server vendors. If you have purchased an encoding system that is not widely supported by client/server vendors, your options are to purchase from the vendor--or provide in-house--the custom API programming needed to create hooks to your files or to select a vendor with very open file support, providing streaming and client playback independent of the encoder card and encoding software used. Always test a system with your encoded files before purchase, however, regardless of vendor claims.

Decoding at the client end is another significant issue, particularly for MPEG-2 service. Many vendors support hardware decoding only and may be very limited in the different cards they support. Again, client/server vendors publish tested decoder cards or should provide this information on request. It is important to know your client population before selecting both the encoding and the client/server system. Do you already have a widely-deployed decoder system that you do not want to replace? Are you supporting a controlled user base, such as a computer lab, or a large heterogeneous user population with varying and unknown operating systems, processing speeds and RAM? Are you supporting a wide range of bandwidths, such as different LAN topologies and dial-up traffic? What percentage of your user population uses Windows, MacOS or UNIX desktops? Which UNIX desktop operating systems must be supported—Linux or a vendor-proprietary OS, such as Solaris, AIX, or IRIX?

The installed user base is a critical criterion for selecting both the encoding system and the client/server system. You can purchase a system only for a managed computer lab but you risk serious user dissatisfaction. Most users won't settle for anything less than full access at their desktops. --And the system will be blamed for performance failures, not the inadequate desktop. If you are serving (and most of us are!) a heterogeneous, somewhat unknown user population, it is best to select flexible systems--encoding systems supporting a range of bandwidths and standards and client/server systems supporting a range of bandwidths, encoding systems, client operating systems, and network protocols. A flexible system will also offer a choice clients--system-proprietary and Web-plug-ins or helper applications.

There are several components of a system or vendor selection process:

Market Survey

The first step in vendor selection is a survey of the market, looking at existing technology, new standards and emerging technologies, customer deployment of current technology, customer satisfaction surveys, the experiences of your colleagues, and, of course, the web pages of all the digital video client/server vendors identified in the market survey. Even in cases where a vendor has already been pre-selected, such as in a statewide digital video initiative, knowing the available technology and offerings of competing vendors is invaluable for working with the pre-selected vendor. In your market survey, concentrate on articles and web sites that survey the market, evaluate existing vendors and technologies and predict future enhancements for digital video technology. Sign up for electronic discussion lists that include users of digital video technology. Ask questions about different vendors and their offerings. Experienced colleagues are probably the best source for information, since they can respond to real-world implementation and management concerns. Sources for information include the following:

- *Colleagues:* Contact any colleagues that have purchased digital video client/server systems. Ask for copies of any of purchase information that your colleague is able to share, such as RFI (Request for Information); RFP (Request for Proposals); purchase order outlining specifications and service requirements; and contract. For government agencies, after bids are awarded, this is generally freely available or, at worst, available through the Freedom of Information Act. Ask for your colleague's experiences with the selection process. Which vendors were eliminated, and why? Ask for experiences with installation and deployment. What level of assistance was provided with set-up and initial troubleshooting? Did the product perform as described? What ongoing issues and concerns are your colleagues experiencing?
- *Vendor References:* Vendors will provide you with references from their clientele. Be sure to check all references from institutions of a similar size and purpose. Verify titles and responsibilities for reference contacts. Is the contact a high-level administrator without sufficient hands-on experience or technical expertise? If so, consider the reference contact a starting point for locating the appropriate person to answer your questions.

- *Websites*: General video information, such as this Website, and also vendor-specific Websites. Look at vendor press releases, generally posted at vendor web sites, for the past year to six months, to get a feel for development patterns in digital video, as well as for each vendor.
- *Discussion Lists*: Subscribe to any electronic discussion lists where digital video users discuss their experiences. Some vendors may host discussion lists for their customers.
- *Conferences*: At conference booths, talk not just to vendors but to other users spending significant time at the booths. Discover their level of expertise and deployment plans and be sure to exchange cards. Current customers often take advantage of conference booths to discuss issues and concerns with vendors. This is a great way to develop a reference list independent of the list supplied by the vendor.
- *Journal Articles*: Articles that survey functionality and customer satisfaction are critical, but don't neglect predictive articles looking at future enhancements. A clear migration path for future technologies is critical for more expensive purchases.

Purchasing Instrument

After a market survey to familiarize yourself with the state-of-the-art for digital video, it is necessary to select the functionalities that are both critical and desirable for your project and to codify those functionalities into a purchasing document, whether an RFI, RFP or a purchase order. If possible, design an instrument that can be sent to a number of vendors.

Be sure that the bid section will result in competitive pricing that can be compared uniformly across vendors. A good practice is to provide a bid sheet with individual line items for each meaningful system component. Meaningful system components vary by project and are best determined by the individual institution, after an extensive market survey. These line items can include the entire system (hardware & software); individual line items for component pieces (streaming server, multicast server, etc.); and line items for services, such as installation, training, and ongoing maintenance. It is critical to request information about warranty and maintenance costs. One often-overlooked pricing differential is warranty period, with some vendors offering three months and others a year or longer. The author has required multi-year bids on ongoing maintenance costs for large-scale purchases, to insure that her organization is able to financially maintain a selected system over time and to insure that vendors do not offset low purchase costs with high maintenance pricing.

For RFPs, although a purchase is implied, be sure to include language that states that the organization you represent is not required to issue a purchase order in response to bids received.

Distribute your purchasing instrument to the widest possible vendor pool. You will probably work closely with your institution's Purchasing Department but do not rely solely on their list of identified vendors. Supplement that list with the vendors you discovered in your market survey.

Your purchasing instrument should require the names and contact information of all customers similar in size and mission to your institution. Do not ask for selected customers, but the complete list of customers meeting your description. A critical component of the selection process is the checking of references. Be sure to ask standard questions of each reference, for comparison purposes, as well as open-ended questions about their experiences. Many vendors may not provide a complete list, even though it is requested. If necessary, ask the contacts provided what institutions or companies they contacted for references, and expand your reference pool in this manner.

Contractual Issues

Depending on your institution and the size and nature of your purchase, no contract may be required, or there may be a purchase contract and a maintenance contract. For expensive projects, where expense includes not just the purchase itself but the staffing and training required for deployment, a purchase contract is a good idea. A purchase contract can provide the following benefits:

Financial and risk protection.

If the contract includes innovations not yet available, the purchase contract can outline staggered payments for scheduled deliverables. The contract can also define financial performance incentives for functionalities that are very new or that do not perform as specified, particularly if you select a vendor for very good reasons in spite of concerns expressed in reference checks about the performance of certain functionalities or problems with ongoing troubleshooting and support. Most vendors have an honest desire to serve customers well, but they are frequently understaffed and focused more on generating new business than on support for existing customers. Financial incentives (also known as financial penalties, when the vendor steps out of the room!) are an effective way to insure service and minimize risk, particularly for very new technologies. Vendors are more likely to agree to financial incentives for performance for large, expensive projects and for projects that will be heavily promoted by the purchasing institution.

For government entities, which of course includes state universities, financial penalties can be tricky but not impossible. Steep reductions in ongoing maintenance costs, free extension of the warranty period, payment in free enhancements, free additional streams, etc. can usually be worked out with your contracts department as well as with the vendor. The goal is to avoid enriching the coffers of your institution's "general fund," which might go toward the purchase of uniforms for the football team, and instead to impose performance penalties that directly compensate your digital video implementation.

Financial protection can and should include price caps for ongoing maintenance and should, if at all possible, lock in prices for enhancements that are part of the purchasing instrument response, and thus the contract, but not yet available for purchase.

Upgrades and Enhancements

A contract is a good place to negotiate for functionalities requested in the purchasing instrument which the vendor is willing to develop but unable to currently supply. Before purchasing a digital video client/server system, be sure to identify, possibly through a non-disclosure agreement, any anticipated enhancements scheduled for release in the next six to fourteen months. If you include any planned enhancements in the purchase contract, be sure to minimize the risk to you contractually. Many vendors offer ongoing maintenance plans that include software-based enhancements. Knowing the vendor's development plans will help you determine the value of a combined maintenance/upgrade plan.

Be careful in a contract to negotiate only for enhancements to current functionalities that would benefit a range of users, such as a MacOS client, for example, and not replacement functionality that would result in the purchase of a non-standard current product. You do not want to risk ongoing problems with new releases and upgrades that will not interoperate with your nonstandard product. If current functionality requires re-working to customize service for your institution, you are probably buying the wrong product. Obviously, very large institutions, such as government entities and consortia, will have better success negotiating functionality upgrades for existing products. If a vendor meets your needs in most areas but lacks one or more key requirements, consider developing a consortial purchase arrangement with other institutions with similar needs.

If you identify a significant enhancement to service that you contract with the vendor to develop, be sure to use the purchase contract or another contract instrument to spell out the specifications and the financial incentives for completion. If your institution's involvement in designing and testing the enhancement will be significant, consider a joint marketing venture, or at least a substantial innovator's discount for the purchase and ongoing maintenance of the enhancement. Make sure all joint venture or pricing arrangements are clearly established in the contract.

SECTION FOUR: DIGITAL VIDEO CLIENT/SERVER STATE OF THE ART: FIRST QUARTER, 1999

Digital video client/server state of the art is discussed in this section in six key areas: open systems design, scalability, functionality, asset management, distance learning application support and purchasing & support. Each area is discussed and illustrated through the product offerings of ten vendors in the digital video client/server market that responded to a Request for Information issued by ViDe in October, 1998. These ten vendors are: Advanced Modular Solutions, Inc./Digital Bitcasting Corp. (joint proposal), CISCO Systems, Concurrent Computer Corporation, CyberStorage Systems, IBM, InfoValue Computing, Inc., Panasonic, Starlight Networks (subsidiary of PictureTel Corporation), SGI, and 3CX Streaming Media Solutions. Information about each vendor was developed from RFI responses and from updated information on corporate web pages. Suggestions are offered for evaluating vendors—those responding to the RFI and therefore discussed in this white paper, as well as vendors that have not responded or have a response in process, such as Progressive Networks (RealVideo), Hewlett-Packard, and Microsoft.

This section should be read with two important caveats: vendor overviews are based on explicit, selected capabilities obtained from RFI responses and vendor web sites. This overview is not intended to be exhaustive, so there will be omissions, for all vendors, for functionality and features. Since this white paper selectively examines digital video issues, some omissions are intentional on the part of the author. There may also be significant omissions due to author or vendor oversight. Finally, there will certainly be omissions, in this volatile field, due to the age of the document. All of the digital video vendors surveyed are actively developing their products. Features not available at time of writing may well be available before this document is revised.

The second important caveat is that vendors are summarized on published capabilities that have not been tested for validity or performance by ViDe. No vendor should be selected without careful testing and evaluation by the purchasing institution. The intent of this document is to acquaint you with the state of the art in the first half of 1999, as a first step in the evaluation and selection process. The most critical steps in the selection process, however, are the careful identification of your unique user needs and the careful evaluation and testing of vendor products to meet those needs.

This section opens with a brief overview of the digital video product offerings of the nine responding vendors:

Advanced Modular Solutions & Digital Bitcasting Corporation:
Offers Progressive Networks' RealSystem G2 for video and audio and MPEG1/MPEG2 using Digital Bitcasting Corp.'s MPEG plug-in for RealSystem G2. Modular's Intel-based server platform, running either NT or LINUX, is provided for video storage and streaming, for a streamed video turnkey system. MPEG4 is in development. Products include an encoding station and non-linear editing station for video asset creation, the Media Archive Server/Digital Library for asset storage, the Stream Server for asset delivery, RealSystem G2 IP multicasting and live broadcasts of MPEG streams. Virage VideoLogger can be bundled for asset management and indexing. The G2 viewer with MPEG plug-in supports Windows 95/98/NT, MacOS (spring, 1999) and UNIX. In addition to client/server digital video streaming and multicast, Digital Bitcasting's 4-in-1 real-time hardware encoder is offered, providing the simultaneous creation of up to four MPEG1 files at varying data rates on a single PC.

CyberStorage Systems:
CyberStorage supports video on demand, IP multicast and live broadcast streams. Components of the CyberStorage video on demand client/server system are the NT-based CBV Server, which maintains video assets in the clip repository and provides asset management, including security and accounting, through its database module. The CBV Transmitter, an NT workstation, accepts and encodes analog live video feeds from cameras, VCRs, etc. for live broadcasting. The CBV Receiver is GUI-based client software for receiving multicast/broadcast transmissions. Video on demand unicast client software is provided by the Clip Viewer. MPEG1 and MPEG2 are supported, as well as any proprietary standards and future standards, such as MPEG4. Proprietary viewers may be needed or the Clip Viewer can be extended to support proprietary formats. The client viewer supports Windows 95/98/NT with MacOS in development.

CISCO Systems:
CISCO provides IP/TV, which supports unicast and IP multicast as well as MBONE broadcasts. MPEG1 and a variety of codecs, including H.261, Indeo, Cinepak, QuickTime, Vxtreme and M-JPEG are supported. MPEG-2 will be supported in IP/TV v. 3.0, scheduled for release in early 1999. MPEG-4 support is also in development. Components of IP/TV include IP/TV Content Manager, which manages configuration and transmission for IP/TV servers and viewers, the IP/TV Server for capture, storage and transmission of video streams, and the IP/TV Viewer. IP/TV can be a software-only solution or bundled with the CISCO IP/TV 3410 Control Server, IP/TV 3420 Broadcast Server and the IP/TV 3430 Archive Server for a turnkey solution. Additional products are SlideCast, providing synchronized slides, such as PowerPoint, with broadcast video, Question Manager, which provides viewer-feedback through questions for immediate response or archive and respond later, and StreamWatch, which provides management information such as viewer demographics and stream quality. IP/TV supports Windows96/98/NT clients and is also compatible with UNIX VIC/VAT for MBONE and the Apple QuickTime streaming extension on Mac.

Concurrent Computer Corporation:
Concurrent offers the MediaHawk Intranet Video System. The component pieces are the MediaHawk Video Server VOD Pack, which includes software modules and libraries to store and deliver video streams, the MediaHawk Player software and the MediaHawk System Administration software providing asset management, statistics and management of the video pumps streaming the video. MediaHawk supports MPEG-1 and MPEG-2. Concurrent is tracking the MPEG-4 standard and considering support for QuickTime. Unicast video on demand and video conferencing are supported by the MediaHawk Intranet Video System. Multicasting will be supported in mid-1999. The viewer client supports Windows 95/98/NT.

IBM
IBM provides unicast and IP multicast streaming through its VideoCharger server, available on UNIX (AIX) and Windows NT. VideoCharger supports MPEG1, MPEG2, QuickTime (through its ActiveMovie implementation), AVI and Bamba (IBM proprietary format). MPEG4 support is in development. Multicast capability is separately priced with a one-time charge. NFS provides file system support on NT servers, while a multimedia file system manages AIX files. IBM Digital Library Version 2 provides integrated digital media management for assets in any format, and the 3466 Network Storage Manager provides video/multimedia archiving for storage and recall of archived media objects. Windows 95/98/NT clients are currently supported, with a MacOS client in development.

InfoValue
InfoValue offers unicast and multicast streaming video through its QuickVideo Suite. QuickVideo Suite supports all video formats as well as all off-the-shelf encoders, decoders and authoring tools. QuickVideo streaming is transparent for file format so expects to support MPEG-4 when available. Components of the suite are: QVOD (Quick Video on Demand) v. 4.0 providing unicast streaming/recording for Windows NT server environments; QVAR (Quick Video Archive), the asset management/video library application; QVMC (QuickVideo Multicast) supporting live broadcast and instant replay; (QVIW) QuickVideo IntraWeb, providing load balancing and management for enterprise-wide Intranet or multiple-site streaming; and QuickVideo Simulcast providing simultaneous real-time encoding of streams for monitoring and after action review. The QuickVideo player supports Windows 95/98/NT clients.

Panasonic
Panasonic Video Network Server is file format independent, supporting MPEG1, MPEG2, and other formats. QuickTime, AVI, Real have been tested and are supported. Unicast streaming is currently supported. IP multicast will be available with VNS 3.1, due to be released in early 1999. Sun's UNIX platform (Solaris) is required for the server. VNS can be bundled with LearningNet, a Java-based authoring application for synchronizing video, graphical aids and course outlines. Client playback operates on UNIX, Windows 95/98/NT and MacOS. A Java media player client is in development.

SGI
SGI offers WebFORCE MediaBase for media streaming, supporting MPEG1, MPEG2, RealVideo, RealAudio and H.263. Plans are in place to support MPEG4. MediaBase runs only on SGI Origin IRIX servers. The Origin 200 server is recommended for smaller installations. Installations wanting 20 streams or less also have the option of the O2 workstation as a server. Unicast, multicast and MBONE streaming is supported. MediaBase currently supports only the Optibase MPEG2 decoder card but comes with a DirectShow filter so that decoding cards supporting DirectShow are also supported. Asset management is provided through a bundled Informix database. Non-bundled Oracle is an option but must be purchased separately. StudioCentral can be bundled for enhanced asset management and indexing. Windows95/98/NT, MacOS, IRIX, Solaris and AIX clients are supported. The Solaris and AIX clients must be licensed separately from a third party partner.

Starlight (Subsidiary of PictureTel):
Starlight's video suite provides unicast and multicast video streaming through the following products: StarWorks, providing unicast streaming and recording; StarCast, providing real-time multicast service; and StarCenter, the file and streaming application manager, which manages RealVideo G2 and NetShow servers in addition to StarWorks and StarCast servers. StarWorks is available for Windows NT and Solaris server platforms. StarLive! provides streaming unicast and multicast with synchronized slides and real-time question and answer capability. StarCenter provides management for any file format supported by StarWorks, StarCast, RealVideo G2 and NetShow. StarWorks and StarCast support MPEG1. StarWorks supports MPEG2; StarCast will support MPEG2 in the next release. StarWorks will support MPEG4 if a player application is provided. StarCenter supports NetShow's MPEG4 implementation. StarWorks provides streaming independent of file format and supports QuickTime, AVI and M-JPEG, among others. ASF and Real can be supported with their native servers through StarCenter. StarCast currently supports only MPEG codecs. Starlight client software supports Windows 95/98/NT.

3CX Streaming Solutions (Subsidiary of IXMICRO)
3CX offers unicast and multicast streaming video through its ixJet Streaming Server and client for unicast, ixJet Live Server and client for real-time IP multicast, as well as the ixJet Streaming SDK and ixJet Live SDK to allow developers to customize and extend video on demand and multicast service. In addition, 3CX offers ixJet Network Video Explorer for on-demand learning and asset management, particularly in the K12 environment, ixJet Video Finder for asset management, indexing and discovery, and ixJet Network Video Presenter for presentation and content creation (timeline driven video, audio, PowerPoint slides, and text annotation). MPEG1, MPEG2, and QuickTime video formats are supported for unicast video on demand and MPEG1 for IP multicast. ixJet Live Server 2.0 will support MPEG2 multicast. MPEG4 support is in development. The viewer client runs on Windows 95/98/NT, MacOS and UNIX

Open Systems Design

Open systems design is a catch-all concept encompassing a variety of standards and processes. Open systems design, based on adherence to standards in networking topologies, file formats supported, file storage and management, data transmission and client reception, is a critical criterion for selecting a system that fits seamlessly with your legacy infrastructure to merge digital video with current services. Open systems support also provides your best guarantee that the VOD system selected will grow with your infrastructure and support future needs.

Another issue to consider when looking for an open standards-based system is the increased amount of collaboration--in the sharing of video assets to create large "virtual" collections; in the sharing of multicast programs, and, over time, the merging of videoconferencing with video on demand as the H.323 videoconferencing over IP standard matures. Increasing demands will be made on digital video systems for distance learning, resource sharing and collaboration across institutions. An open design may not be a critical issue at time of installation but will grow rapidly in importance. You do not want to commit to a proprietary solution, and encode numerous assets, only to find the system will not scale beyond your department or institution.

A completely open system will transparently support digital video files in any format, a wide range of encoding and decoding hardware and software, to a range of clients, including Windows, Mac and UNIX. Service will be provided by Windows and UNIX servers, utilizing a variety of TCP/IP and UDP/IP transport protocols, including RTP, RTCP, RTSP and IGMP. An open system will have complete web capability, including web plug-in and helper-app client software, remote server and file management via the web, web-based search engine support, and as well as http file access and security.

A truly open system should interface transparently with the systems of other vendors so that two or more institutions could collaborate to create “virtual” collections and services to support user needs. Imagine, for example, users at two institutions, with different digital video client/server systems, transparently sharing a video newsreel collection on life in the U.S. during World War II at one institution and documentary war footage at the other institution. In addition, history lectures from both institutions are transparently viewed via IP multicast by users at either institution.

Full interoperability between disparate systems is a future enhancement that ViDe hopes to explore in the near future with selected vendor partners. Currently, this interoperability is not available, or at least not actively developed and tested, by digital video vendors.

Components of open systems design are discussed below:

Open Networking
Simultaneous unicast streams require a large amount of bandwidth. As your needs grow, it is important to be able to migrate your services to a higher-bandwidth network. Also, future applications involving full MPEG-4 functionality--accompanying text, 3-Dimensional images, object indexing and manipulation, etc. will require more bandwidth, even if the stream number stays the same. An open system will support every flavor of Ethernet, including switched-10, 100BaseT ("Fast Ethernet"), Gigabit Ethernet, as well as ATM (ATM with classical IP (CLIP), ATM with LAN emulation (LANE), and native ATM with Quality of Service (QoS). Depending on your existing topology, or the need to collaborate with another institution or department with a different topology, Token Ring support may also be important to you. Transport protocol support, including TCP/IP, UDP and RTP/RTCP is also a critical requirement for managing network traffic and stream quality. For dial up access, cable and ADSL support can be important considerations.

The ten vendors participating in the RFI are fairly strong in open networking support. Most responded that they support Ethernet and any TCP/IP based network topology. IBM, InfoValue, SGI and Starlight explicitly support FDDI. IBM and InfoValue explicitly support Token Ring. CISCO, Concurrent Computer Corporation, IBM, InfoValue, SGI and 3CX support ATM, with InfoValue, SGI and CISCO explicitly supporting LANE, CLIP, and native ATM over QoS. Panasonic states support for "segmented Ethernet 10BaseT protocol. Co-exists with other popular networks." CyberStorage states support for "All major LAN & WAN interfaces." Advanced Modular Solutions & Digital Bitcasting Corporation support Ethernet and IP-based network topologies.

When selecting a vendor, verify network topologies and bandwidths supported, as well as ADSL and cable modem interfaces, for remote access.

When selecting a system, verify with the vendor which delivery protocols are supported. Protocol support should include both TCP/IP and UDP/IP protocols, including RTP, RTCP and RASP. RSVP, for resource reservation, is a good protocol to support, for long-term scalability. CISCO and IBM explicitly support RTP, RTCP, RASP, and RSVP. No information is currently available for Concurrent Computer Corporation, CyberStorage, InfoValue, Panasonic, SGI or Starlight. Advanced Modular Solutions/Digital Bitcasting support TCP, UDP, RTSP and PNA protocols. 3CX notes support, more generically, for TCP/IP and UDP/IP protocols. IGMP (IP Multicasting) should be supported for multicasting. Advanced Modular Solutions/Digital Bitcasting, CISCO, IBM, InfoValue, SGI, Starlight, and 3CX explicitly note support for IGMP. Concurrent Computer Corporation and Panasonic do not currently offer IP multicasting service.

Server Platform Support
Support for off-the-shelf third-party servers, rather than vendor-proprietary, and for NT and UNIX operating systems offer the most purchase and deployment flexibility. However, open networking support (discussed above) and scalable server configurations (discussed later) are more critical issues, unless you have a need for a specific server operating system or if low cost is a driving factor. Windows NT servers do not scale as well as UNIX servers, but they will be less expensive than UNIX and can be supplied off-the-shelf by a third-party server vendor, giving you more options for comparison shopping.

CISCO, CyberStorage, InfoValue, and 3CX require NT servers. Starlight supports both Windows NT and Solaris for its StarWorks VOD server but requires Windows NT for its StarCenter management server, StarCast multicast server, and StarLive distance learning/streaming media application server. SGI requires its proprietary UNIX (IRIX) server platform--the Origin series. For very small (1-20 stream) installations, the IRIX O2 workstation is an alternative. Panasonic requires a Sun Solaris server. Concurrent requires a server with its the PowerMAX OS proprietary UNIX operating system. Advanced Modular Solutions/Digital Bitcasting offers NT and LINUX-based stream servers. IBM offers both it's proprietary UNIX (AIX) and Windows NT server solutions.

File Format Support
One of the most important considerations is file format support. All ten vendors support the MPEG standard, but implementations of these standards vary widely. Question the vendor closely about MPEG implementations, particularly in the following areas.

MPEG1 and MPEG2 Unicast:
Nine vendors responding to the RFI offer MPEG1 and MPEG2 support, but not necessarily for all encoding and decoding systems. CISCO offers MPEG1 Unicast, with MPEG2 Unicast available with version 3.

It is critical to verify support for the encoding system already in use to create video files and for legacy hardware and software based decoders on client workstations. I cannot over-stress the importance of this issue! If you have not purchased an encoding system, the author recommends selecting the client/server system first, or selecting the two in tandem. All vendors should test encoding systems and decoder cards for compatibility with their client/server streaming software.

InfoValue states that it supports all available off-the-shelf hardware encoders and decoders. CyberStorage states that it supports any type of digital content and offers automatic support for any hardware decoder.

MPEG2 encoding and decoding support is particularly critical. Vendors may require hardware decoders and may only work with specific models. Some vendors only support MPEG2 streaming at the lower bandwidth spectrum, which is an important consideration for encoding. Bandwidth range--a critical issue!--is discussed in a later section.

Of course, you will want to verify support for your encoding and decoding systems of choice, for any vendor. The author would recommend putting full support for your selected encoding system(s) and hardware or software decoding system(s), including the required bandwidth range and/or encoding level, in the purchase order and/or purchase contract as a condition of sale, so that you are assured of complete compatibility between the files you create, the system storing and serving those files, and the client desktops receiving and playing back the files. Vendors support the constrained bit stream 1.5 Mbps for MPEG1, but a wide range of bit rates for MPEG2. Bitrates (single stream and aggregate) are discussed in the *Scalability* section below.

MPEG1 and MPEG2 multicast
Concurrent Computer Corporation and Panasonic do not currently support IP multicasting. Panasonic will offer IP multicasting in VNS release 3.1, due in early 1999. CISCO, Starlight, and 3CX currently offer MPEG1 multicasting but will support MPEG2 multicasting in the next release. All other vendors responding support multicasting, specifically IGMP (Internet Group Management Protocol) for IP multicasting. Those vendors did not specify MPEG2 multicasting.

Multicast streams may be handled differently from unicast streams. Verify whether MPEG1 and MPEG2 streams are both supported, as well as the bitrate range supported for multicast. Verify with each vendor particularly whether MPEG2 multicasting is supported, as well as the transmission bitrate. It is not uncommon for a lower MPEG2 bitrate to be supported for multicasting than for unicast. If high-bandwidth multicasting, in applications such as digital medicine or modeling, is critical to you, be sure that the vendor can support MPEG2 multicasting now, or that this service is in development for a contractual commitment. Also ask whether video and audio are transported as separate streams and, if yes, how synchronization of the streams is assured.

Vendors generally support both live and stored multicast broadcasts. Some vendors bundle encoders and recording software so that a live analog feed can be encoded on the fly. An important question to ask is whether a specific encoder/recorder is required for live broadcasting. If so, you will want to insure that files are produced in a standard (MPEG1, MPEG2) medium to high-bandwidth format to create and store assets of a consistent quality with those you already own.

Low-Bandwidth File Format Support
Systems that offer completely open file streaming support will be indifferent to encoding format for storage, basic asset management and streaming. Starlight, Panasonic, InfoValue and CyberStorage all state transparent unicast streaming support for any file format. Advanced Modular Solutions/Digital Bitcasting state that they offer "open file streaming support" with the G2 Real Server. Open file streaming support can be easily verified during a vendor demo by providing files in a range of low-bandwidth formats, since encoders are either free or very inexpensive for most low-bandwidth formats.

Low-bandwidth videos in different formats and bandwidths can also be borrowed extensively from the web, with permission of the page owners, for demonstration purposes. It is always a good idea to test streaming and client functionality with your own files, as well as the canned files provided by the vendor. Before beginning the demo phase of your evaluation process, gather a number of video files (high-action, high-color, talking head) and audio files (range of speaker voices--high pitched, low-pitched, etc., range of music types, including singing and instrumental only) for testing across vendor platforms.

While MPEG-4 has been adopted as a standard, it is not yet available except in alpha and beta test, for encoding, storage, streaming and playback. Proprietary low-bandwidth formats must be supported for Internet streaming outside the building or institution LAN. All vendors responding support one or more low-bandwidth formats, in addition to MPEG1 and MPEG2.

CISCO states that IP/TV supports a variety of codecs, including H.261, Indeo, Cinepak, Vxtreme and M-JPEG but notes "in some cases, specific hardware encoder or decoder cards may be required." IBM 's VideoCharger supports MPEG1, MPEG2, QuickTime (through its ActiveMovie implementation), AVI and Bamba (IBM proprietary format). 3CX supports QuickTime. SGI supports RealVideo, RealAudio and H.263. Advanced Modular Solutions/Digital Bitcasting provides the Progressive Networks' RealSystem G2 for video and audio with MPEG1 and MPEG2 as a plug-in application to the G2.

Which proprietary low-bandwidth formats are best? This is mostly an application-specific decision. In the author's opinion, RealVideo and Audio are widely available, generously supported with very good free clients, free encoders, free plug-in applications, free server software (20 streams at date of writing), etc. File quality and streaming are excellent, as is support for the SMIL (synchronized multimedia integration language) standard for incorporating digitized text for audio file transcripts and distance learning. Other new features include RealFlash, RealMedia's integration of Macromedia Flash into its authoring product. The RealSystem seems determined to maintain its enormous presence on the web. If plug-in integration for streaming MPEG-4 is offered, and the current bandwidth streaming limitation of 900KB overcome, RealMedia will remain an excellent low-bandwidth solution.

QuickTime's file format provides the wrapper for MPEG-4. The file format is excellent for download and play, but until recently, streaming capability was lacking. QuickTime 4.0, which was released in April 1999, incorporates streaming functionality and protocol support. In addition, QuickTime 4.0 provides built-in support for digital video, including MiniDV, DVCPro, and DVCam camcorder. QuickTime 4.0 will include source code, allowing great flexibility for custom APIs. The compatibility with MPEG4, while not at the object recognition and manipulation level, is a real plus.

At the present time, the author would look for both RealMedia G2 and QuickTime 4.0 file formats to be supported. At a minimum, insure that at least one of the two--whichever interests you the most for your applications--is supported. Mature MPEG-4 solutions are probably a year away, and a year is a long time in the fast-paced world of the Web. RealMedia G2 and QuickTime 4.0 offer significant enhancements over previous software versions. Be sure that the video client/server system supports the latest release for your proprietary low-bandwidth format, or has firm plans to do so

Client Support

Your client base is the end user PC platforms that will play your streaming video files or multicast transmissions. The client base can rarely be completely identified in advance and is therefore a critical concern for most installations. Your client base includes workstations you directly support, in the building where the server resides, customers accessing videos remotely from home or office, as well as any authorized visitor to your web site. Be careful about the assumptions you make concerning your client base. The best way to discover just how many MacOS clients you support is to offer a service with a Windows-only client!

All responding vendors provide Windows players. Support for MacOS and UNIX is another story. Advanced Modular Solutions/Digital Bitcasting uses the RealMedia G2 client software, supporting UNIX, Windows 95/98/NT and, in spring 1999, MacOS. CISCO’s IP/TV supports Windows 95/98/NT and can interoperate with the MBONE VIC/VAT client on UNIX platforms as well as the QuickTime streaming extension on the Mac. Concurrent’s MediaHawk Player runs under Windows 95/98/NT, and under any application that uses Microsoft DirectShow 5.0, including ActiveMovie.

CyberStorage supports Windows 95/98/NT with MacOS in development. The CyberStorage response notes that UNIX and MacOS can use NFS protocol to access the server and use native media players rather than the CyberStorage client. IBM supports Windows 95/98/NT and has a MacOS client under development.

InfoValue’s QuickVideo on Demand player runs on Windows 95/98/NT clients. Panasonic offers both a Solstice NFS client and a Java media player, utilizing the browser-embedded ActiveMovie player. Client playback operates on UNIX, Windows 95/98/NT and MacOS.

SGI supports Windows95/98/NT, MacOS, and IRIX. SGI also provides client support for Solaris and AIX, which must be licensed from a third-party developer, RABA. The ixJet Streaming and Live clients from 3CX runs on Windows 95/98/NT, MacOS, and UNIX. An SDK for the Streaming and the Live Clients provides Java-based API development including a Java applet for a cross-platform client plug-in.

Verify client performance on your existing UNIX and MacOS client desktops. It is critical to insure that the vendor's players work with the versions and flavors of OS software that you support. UNIX clients may not run on Linux, a very popular UNIX desktop OS, for example. Compare the Windows clients with MacOS and UNIX clients for functionality. Are all features, such as fast forward, rewind, freeze frame, bookmark, resize, etc. supported? If not, are there plans to add missing features? If you must use a native or third-party client rather than the vendor's client, look for tradeoffs in appearance, functionality and streaming performance. Vendor players will include proprietary communication with the server for managing the stream, which could noticeably affect performance.

Web Support
The Web has become a kind of meta-operating system, providing connectivity, database and document management and support, as well as a common interface for computer-based communications of all kinds. Web integration and support is a very important component of any digital video client/server system. Client/server systems use the Web and its supporting protocols for access to digital video files, but also for asset management, indexing, security and remote server administration. At a minimum, a VOD client/server system should support http access—both for asset streaming and remote management for asset loading, backup and file transfer among servers. In addition, Web-based security and client playback through a Web page should be supported.

3CX’s ixJet Streaming Server and Client and ixJet Live Server and Client include an SDK with ActiveX control APIs for Windows applications and a Java package for cross-platform applications. The SDK includes HTML embedding for Web pages and for creating a client plug-in using either ActiveX control or Java applet. IxJet Network Video Presenter, the Web-based authoring and presentation tool, supports the Secure Socket Layer for controlling access to distance learning applications through user authentication and data encryption. InfoValue provides an API that can be used for Web-based applications and is compatible with a wide range of development tools, including Visual Basic, Visual C++, Macromedia Director, among others, including database tools. InfoValue’s API is designed to be open and transparent to the application development tool, with the stated goal that “the best API is no API.”

Advanced Modular Solutions/Digital Bitcasting provides the Web capabilities of the RealMedia G2 client/server software, including SMIL and Flash authoring and playback capabilities in a Web environment. In addition, authentication is based on the HTTP standard, RFC 2069. CISCO’s IP/TV includes a Web plug-in for launching and configuring on a Web site as well as a web button on the viewer to launch a predefined Web page. SGI’s WebFORCE MediaBase also uses a browser plug-in that can be embedded in Netscape Navigator or Internet Explorer browsers.

Panasonic provides a Java viewer for Web-based access, as well as remote administration of the server via the Web. IBM’s DB2 Digital Library VideoCharger Player provides an external helper view application as well as a browser plug-in viewer for embedding the viewer in a browser window with a subset of the standalone viewer controls. IBM also provides a video-on-demand Web page for searching and selecting videos for playback.

Other Open System Components

In addition to the open system components listed above, it is important to look at several other network components when evaluating vendors:

Security. Does the vendor support http security features? Will the system support kerberos? Will the vendor's own security system interoperate with the institution's security system? Security features to investigate include authentication (access to service), authorization (access to specific files within the service) and encryption (security at the file level). The levels of security to be supported depend on your needs. You may want encryption for multicast distance learning applications where a fee for "attendance" is required, for example.

Database. What database is used to provide file management and administration? Is the database ODBC-compliant? --CORBA compliant? Many vendors provide a database management system such as Access, Informix, etc. but also provide an API for Oracle.

Storage. A digital video client/server system should support a wide range of storage and HSM (hierarchical storage management) options, including JBOD ("just a bunch of disks") arrays, RAID, optical storage (CD-ROM and DVD), FC-AL (Fiber-Channel Arbitration Loop) as well as storage methods such as robotic manipulation and jukeboxes.

Application Standards Support. Most applications are governed by a suite of standards specific to the application, such as SMIL (synchronized multimedia integration language) for multimedia presentations. The Real G2 client supports the SMIL standard. Asset cataloging and indexing includes numerous standards for creating catalog, or metadata, records, for searching metadata records and retrieving related assets. IBM supports the Dublin Core metadata standard, a robust, general purpose metadata standard for cataloging Web information, and also the Information Retrieval (Z39.50): Application Service and Protocol Specification, a client/server network protocol that enables a client to access any Z39.50 enabled server (one or many) to search databases, retrieve records and organize and display the retrieved records. Z39.50 is not yet a mature standard but has the capability to greatly expand access to information at the database record level. 3CX offers metadata support compatible with the IMS (instructional management systems) Meta-Data standard in its ixJet Network Video Explorer component. IMS Meta-Data is a cataloging standard proposed and supported jointly by Educom and the National Learning Infrastructure Initiative (NLII) and is intended to standardize access to courseware and distance learning tools.

Scalability

Scalability is the ability to provide expanded service, as use increases and needs change, beyond the requirements of the initial installation, without replacement of major system components. Scalability, in this white paper, includes

- support for more simultaneous streams;
- support for more video files;
- file service to additional locations;
- more bandwidth, even if file and stream quantities do not change. An example would be the decision to move from MPEG1 to MPEG2 for the file format.

What are the elements of a scalable system? The most critical is support for an efficient distributed server architecture (or, alternatively, a centrally-located server farm), where files can be stored on many different servers and selected for streaming according to availability, user authorization or location. File distribution across servers requires centralized file management for efficient use of bandwidth. Centralized file management should provide load balancing among servers, at a minimum. The best file management systems will push assets to servers closest to the heaviest end node use and will select among duplicate assets according to availability, bandwidth requirements and network path. All servers should support remote management via the Web.

A distributed server architecture should also insure high availability, providing hot failover to a duplicate asset on another server, without interrupting the stream, so that server failure is imperceptible to the user. Windows NT and UNIX both support server clustering. It is critical to verify, however, that clustering and hot failover occur at the application (e.g. video on demand and video multicast) level and not just at the native file system level, to insure transparent service to your users.

The distributed server architecture with centralized file management provides a way to grow your system without significant management overhead. There is considerable up-front expense involved in a distributed server system. The initial server purchased should be scalable to support at least two years of need without upgrade. A scalable server should support high bandwidth streaming--for a single stream and for the aggregate bandwidth of multiple streams. The initial server should support 100 to 300 streams. High speed LAN protocols and high-density storage devices should be supported, as well as failover clustering for high availability.

All ten systems that are discussed in this paper support multiple high-density and high-speed storage options, as well as high-speed network protocols. Most vendors offer centralized file management options for video network expansion. IBM's RS/6000 server family with AIX uses data pump architecture where data pumps are added on additional servers but all servers and data pumps operate under the management of a single control server with a single system image. The RS6000 SP system allows a single node to serve as the control server while additional nodes provide stream delivery. The IBM RS6000 model F50 can deliver 200 streams at 1.5 Mbps, which scales to 1,000 through the addition of more servers. IBM's Digital Library product separates object management from storage and streaming management. IBM also offers the Network Storage Manager for managing high-capacity storage and disaster recovery.

InfoValue's QuickVideo IntraWeb (QVIW) product provides file distribution across servers utilizing "central" push to local servers and distributed caching to manage both ends of the streaming process. InfoValue's QuickVideo on Demand (QVOD) server supports 300+ MPEG1 streams from a single Windows NT server.

CISCO's IP/TV provides replication and distribution of assets among servers and performs load balancing by assigning requests to the most appropriate server, based on content availability, server availability, load balancing, authentication, bandwidth management criteria, etc. The Content Manager controls assets by querying each server, updating program information and logging information to a database. Content Manager is Java-based and can be used remotely via a standard Web browser. StreamWatch, a separate product, provides information on bandwidth utilization by audio and video streams, viewer usage patterns and playback quality at the client. Each StreamWatch program can monitor five simultaneous streams. Multiple StreamWatch copies can be used and the statistics aggregated to a single database for output. CISCO's IP/TV 3430 Archive Server provides more than 35 Mbps of streaming performance--approximately 25 MPEG-1 streams, and can be clustered to scale to 1,000 streams.

SGI's WebFORCE MediaBase servers can be configured with IRIS FailSafe software, to automatically failover in case of server failure in a clustered environment. Failover can be "active/standby" where one server remains in standby until needed or "dual active" where each server works until required for failover. The Origin server platform is field upgradeable with a "pay as you grow" model for adding node cards. Each node card contains two processors and can handle 150-200 1.5 Mbps streams. Origin servers range from the Origin 200 (scalable to 300 streams) to the Origin 2000, which can support 64 node cards (9600-12,800 streams). CISCO's Local Director product can be purchased to manage front-end traffic and load balancing on a server cluster with distributed assets.

Starlight's StarWorks VOD server supports 100 streams at 1.5 Mbps with a Pentium II-based server with dual PCI buses. Servers may be clustered together to aggregate to 1000 streams or more, with centralized file management provided by the StarCenter product. In development is the ability to share files between clusters, each independently managed by a StarCenter server, which has potential for campuswide and consortial implementations. StarCenter can be provided as a failover service through a commercially available checkpointing and mirroring product. StarCenter supports media servers from Starlight, Microsoft and Real Networks. StarCenter distributes requests for assets according to client capability, authorization and resource availability, providing the highest-quality version of an asset for which bandwidth is available.

Panasonic's Video Network Server, running on a Sun Solaris platform, supports a maximum file size of 1 TB of data and a maximum file system size of 1 TB. Four multimedia file systems, which can each be optimized to a specific codec, can be supported within the VNS software. Content can be remotely managed. Content replicated at varying bandwidths is made available at the highest bandwidth supported by the client, based on the connection type registered by the client at log-in. Sun offers several cluster topologies for providing failover capability and for adding video streams.

CyberStorage does not currently provide clustering for video on demand with all the features of the CBV software. Remote sharing and file maintenance is supported. CyberStorage supports up to 300 users on a single server. The CBV server can be multi-homed with multiple physical network interfaces and is designed to provide load balancing across its configured network interfaces The vendor states that 100s to 1000s of users can be supported through shared computing and storage. CyberStorage offers integrated high-speed storage solutions with its Ultra Fast storage system providing a sustained transfer rate of more than 115 MB/sec.

Concurrent Computer Corporation sells a turnkey solution incorporating its MediaHawk Video Server. Concurrent provides servers that support from 10s to 1000s of MPEG1 and/or MPEG2 streams, up to 12 Mbps. MediaHawk servers are field-upgradeable with hot-swappable disks and can be configured for no single point of failure. File management is provided by the MediaHawk System Administration Software, which holds and manages video files, provides server configuration, system performance monitoring and diagnostic features.

Advanced Modular Solutions/Digital Bitcasting Corporation's Modular Video Server supports sharing of media storage devices by all servers. Multi-terabyte storage is supported. Servers are rackmounted, modular components that can be configured for scalability and redundancy.

The ixJet Streaming Server from 3CX delivers up to 100 MPEG1 video on demand streams per server. 3CX's multiple server solution (in development in November, 1998 for RFI response) can scale to more than 800 streams using eight PCs connected to an FC-AL storage network with a throughput rate of 100 Mbps. A server manager program provides for the creation, export, download and deletion of video assets. In version 2.0, remote management via the Web will be provided.

Bandwidth Supported

Standard information was not received from all vendors indicating bandwidth ranges supported for each file format. MPEG-1 bandwidth ranges should include 1.5 - 3 Mbps, while MPEG-2 should support, at a minimum, 2 Mbps through at least 10 Mbps (broadcast quality but less than HDTV quality). Depending on your application needs, you may want to require MPEG-2 per-stream bandwidths in the HDTV range (12-20 Mbps). Vendors should use 1.5 Mbps to indicate number of simultaneous MPEG-2 streams supported, for aggregate bandwidth. Advanced Modular Solutions/Digital Bitcasting Corporation states they will support on-demand or live MPEG streams at bandwidths ranging from 100 KBPS to 15 Mbps. However, for live and simulated live streams only, RealServer G2 currently imposes a ceiling of 900 Kbps. Digital Bitcasting states that this "maximum bit rate will be increased in the coming months." This ceiling does not apply to video on demand streams.

IP Multicast Streams

Bandwidth ranges for IP multicast will vary from the bandwidth ranges supported for video on demand streams, so each range should be verified. Also verify the maximum number of clients per multicast stream, whether unlimited or a specified number, such as 900.

Asset Management

Asset management includes both digital video files and users, for authentication and accounting for "pay-per-view" applications. Basic file management, particularly for assets distributed across servers, involves loading, replicating, distributing and deleting files. File management should be available remotely via a standard Web browser. Several features of asset management, beyond basic file control, include:

- Management of title copies in different formats and bandwidths. This management should include streaming by the server to the client of the selected asset copy based on factors such as bandwidth at the client, bandwidth availability on the network and best available copy availability.

- Security. Security should include support for authorization by password, IP address or other scheme. File security should support encryption.

- Video indexing is a developing area and can include structured metadata indexing, keyword and natural language searching, relational table indexes, storyboard creation and indexing and visual object recognition.

The capabilities of each vendor are discussed below:

CyberStorage Systems' Media Server provides an SQL database to define users and access to content. The Media Server will not automatically identify duplicate assets transcoded to different formats or bandwidths but does support storing, cataloging and separating duplicates so users can only retrieve the appropriate asset for their system. Tracking and billing functionality is provided for "pay-per-view" applications. Tracking is accomplished by requiring login to the CPV server before any assets can be requested and streamed. Assets may be categorized through customizable catalog fields.

3CX supports Windows NT role-based security and the secure socket layer (SSL) for data integrity over the Web. Advanced file management and indexing are provided with the Network Video Explorer, a multimedia applications product that includes Microsoft's Jet database engine and a relational database for storing metadata that "follows closely" the IMS Meta-Data standard. Network Video Finder is a web-based video context search tool that includes a category builder for three-tier information retrieval. The administrator can design a custom information structure and complete the entry form to create searchable context for each video asset. A built-in video search engine indexes the context database and retrieves video assets from search queries.

IBM's VideoCharger VOD server and Digital Library management system support the transcoding of digital video titles into multiple formats and bandwidths but do not support automatic streaming of different formats/bandwidths based on established parameters. Instead, all instances of a title would be grouped into a title folder, which could be discovered through a parametric search, and the appropriate copy selected by the user for streaming. VideoCharger includes a kerberos ticket for each video asset as part of the metadata cataloging to set playback limits on each video asset. Digital Library supports a variety of digital objects, including still images, video, audio and text. Digital Library provides a variety of indexing schemes, including Dublin Core or a user-designed database model. Z39.50 can be employed for searching and sharing assets across institutions. QBIC (Query by Image Content) is provided for object component recognition (shape, color, texture, etc.) within still images. A similar object recognition product for video is in development. IBM's digital watermarking technology for images is also provided Digital Library can manage other streaming servers with an API to develop hooks into the file structure. IBM has prepared an API for RealServer support.

Panasonic provides management for media, categories and users. The user database provides the system security features through its logon function. No information was provided on cataloging and indexing capabilities.

Starlight's Management product, StarCenter, relies on Windows NT security but is extensible to other security mechanisms. StarCenter manages multiple encoding instances of a title according to three factors: client capability, administrative policies at the server and current resource availability. Based on these factors, the best available copy is streamed to the requester. StarCenter is interoperable with third party indexing systems and plans to develop a tight integration with a third party indexing product as a future development. StarCenter supports streaming servers from Starlight (StarWorks), Microsoft and Real Networks. StarCenter provides usage tracking, file replication, authoring assistance and media clip management.

SGI's WebFORCE MediaBase supports security and authentication for URLs and provides an API (AAA) for enhancing security. MediaBase is sold with an Informix database, which stores metadata, including user-defined keywords and comments about each asset. The asset itself is stored in a real-time file system, separate from the relational database providing query and retrieval. Oracle can be substituted for Informix. Database set-up and management is hidden for Informix but would be necessary for Oracle implementation. Multiple formats of the same title are stored as a composite title. MediaBase uses a feature called Client Capability Based Streaming to query the MediaBase client, determine its bandwidth capability and stream the appropriate format. StudioCentral, a bundled third-party asset management environment, can create a complete media repository, including customized data types, metadata formats and storage mechanisms. StudioCentral is database independent; the StudioCentral Developer's Kit provides database keys and extended type facilities for Oracle8 and Informix Dynamic Server.

CISCO IP/TV's Content Manager offers password protection, encryption and file management through a built-in database, including keyword indexing. Keyword queries can be made through the IP/TV client to identify media assets. CISCO's StreamWatch provides viewer and asset information for up to five simultaneous streams. Viewer information includes identification of the viewer (physical location, network address, email address, etc.) and time viewer joined and left a given program. StreamWatch provides quality information about the data stream itself, including jitter, number of packets received and lost, and amount of bandwidth utilized. StreamWatch can trigger bill-back accounting and can be reported into a database for spreadsheet and report generation.

Advanced Modular Solutions/Digital Bitcasting supports player-based or user-based authentication, using http security protocols. The Virage VideoLogger is bundled with the Real Networks G2 server and client for asset management. Virage VideoLogger uses advanced image analysis to create content-based segments based on scene changes, camera changes, etc. Distinct keyframes are extracted to create a digital storyboard for visual content indexing. The index thus developed can be output to a variety of formats, including a flat file or database management system.

InfoValue's QuickVideo Archive (QVAR) organizes media assets into databases for interactive browsing and playback. InfoValue states that it seamlessly supports all third-party applications, including security, indexing and retrieval systems.

Concurrent's MediaHawk server requires administrators of the system to log in and provide a username and password for access to video assets but would also support any UNIX-aware security mechanism, such as kerberos. Currently the following index fields are supported for each asset: title, length (minutes and seconds), number of frames, frame rate per second, and asset file size.

Distance Learning Application Support

Transparently, or with minimal HTML programming, a video client/server system should support any application, whether a PowerPoint presentation, a multimedia authoring package, such as Macromedia's Director, or distance learning courseware package. If your institution is using a specific courseware package or officeware suite, verify that API's exist for those software packages or that integration is a simple matter. Verify that the vendor offers an SDK (software developers kit) for integration with any application not currently supported by an existing API. The SDK should be Java-based, for crossplatform applications.

Several digital video system vendors offer distance learning applications as separately-priced modules. Starlight offers the StarLive product, providing synchronized slides, streaming video and online chat for audience voting and question/answer sessions. Advanced Modular Solutions/Digital Bitcasting provide, through the Real Server G2, support for Real Producer Plus and Real Producer Pro, presentation authoring software that combines SMIL and HTML templates for creating synchronized multimedia presentations. CISCO's IP/TV offers SlideCast and QuestionManager for presentations with synchronized slides and online chat sessions for questions and answers. Panasonic's Video Network Server offers the bundled LearningNet from Softcom that integrates audio, video, program outlines, graphics and text into a synchronized presentation.

3CX offers two products: Network Video Explorer and Network Video Presenter. Network Video Explorer is an on-demand learning tool designed to organize, manage and share knowledge resources. Network Video Explorer supports asset streaming, hypermedia browsing and end-user configuration and composition for customized applications. Network Video Presenter streams MPEG, Real and NetShow video and audio with synchronized slides and video annotation in a timeline authoring environment.

Purchasing & Support

Most digital video vendors offer separately priced products for each component of the digital video client/server application, including unicast streaming, multicast streaming, asset management, etc. These products are of course interoperable. The institution must decide which products should be purchased initially. Each product has its own development and upgrade release schedule, but interoperability among products should not be affected. If you are purchasing some products now and others later, it is best to verify continued compatibility, however.

Some vendors offer turnkey hardware/software solutions. In most cases, the bundled hardware is optional. If hardware is not optional, be certain that the system will meet your needs, including any consortial or regional file sharing, for at least five years into the future. A portable software solution is generally a safer purchase in the volatile world of digital video, even if you elect to purchase the vendor's bundled hardware platform.

Vendors offer a variety of purchasing options. The most critical issue for VOD pricing is whether the unicast or multicast client/server software is "per seat" or "per stream." Per seat (AKA "per client") pricing requires that each client be accounted for. If you are serving videos to an unknown audience--perhaps via open multicast or a widely available Web page--it is critical that you are able to freely distribute the client. If, on the other hand, you are serving videos to a controlled user population, such as a movies-on-demand service in campus dormitories, per seat pricing may not be an issue.

Vendors offer a range of support options, including upgrade support providing access to new releases and product upgrades at no extra cost (or a greatly reduced cost). Hardware and software support should include a telephone help line, problem escalation and a published response time for problems.

Disabled Accessibility Support

An emerging area for digital video and audio is reliable, transparent access for disabled users. Vendors were not surveyed in this area in the RFI, but ViDe considers accessibility to be a critical development area for the coming year. One of the transformative aspects of technology is its ability to provide greater access to information for people with disabilities. The blind and visually-impaired, for example, benefit from the addition of audio descriptions to video signals. Audio descriptions, which are becoming more common for PBS and cable programming, describe screen action, body language and other visual cues. Examples of audio description tracks attached to digital video can be found at the National Center for Accessible Media (NCAM) Website. (http:// MoviePlayer for Macintosh (version 2.1 or higher) is currently the only client that allows users to toggle the audio description track on and off.

Digital video streamed at low-bandwidths usually renders finger spelling and facial expressions confusing or unintelligible to the deaf and hearing-impaired. This problem can be addressed through the addition of captioning tracks to video signals. MoviePlayer for Macintosh (version 2.1 or higher) is again the only client that currently permits closed captioning (allowing users to toggle the captioning track on and off); other clients, such as RealPlayer, support open captioning (where users cannot turn captioning off). Once again, examples are available on the NCAM Website (http://www.wgbh.org/wgbh/pages/ncam/). Captioning functionality is included in MPEG4.

Captioning offers a number of additional benefits for both hearing-impaired and hearing users. The transcripts created by captioning can help users determine whether a video is relevant to their needs, and can also serve as a partial substitute for the video itself. Some video clients include a search feature that allows users to scan the text track for a keyword or phrase, enabling users to jump to a particular point in the video. The combination of captioning along with audio and video tracks also has potential as a powerful educational tool for teaching both children and adults how to read their native language or foreign languages. Multiple captioning tracks can provide simultaneous translations for assets intended for an international audience.

All of these benefits, however, depend on the development of clients that seamlessly integrate additional tracks, as well as the development of authoring tools that make adding such tracks as transparent as possible for those with limited technical skills.

Providing additional audio description and captioning tracks for real-time streaming media remains a challenge, both in terms of technology and cost, but the benefits for all users are significant.

Still have questions about digital video on demand??? Email the author, Grace Agnew: grace.agnew@library.gatech.edu

New Directions for Video on Demand

Digital video is at an exciting crossroad: the technology is maturing to the point that users are demanding the same robust capabilities--for access and delivery--that electronic text systems provide. At the same time, new developments such as MPEG-4 offer the opportunity to easily create synchronized interactive multimedia objects that are experienced rather than viewed.

The Video Development Initiative (ViDe) began as a collaborative effort among five institutions--University of North Carolina, Chapel Hill, Georgia Institute of Technology, the University of Tennessee, Knoxville, North Carolina State University, and NYSERNet--to promote the use of digital video among academic institutions through knowledge sharing, collaborative projects and product development. ViDe will be working with digital video vendors to improve robustness and ease of use as well as to expand digital video capabilities. ViDe's goal is to integrate digital audio and video into every part of the university learning process: the classroom, the research paper, the collaborative project, even the frantic last-minute review for a final exam.

ViDe provides training workshops, Web-based information, such as this white paper and the Video Conferencing Cookbook consulting and development support for digital video projects and partnering with vendors to develop enhancements and solve continuing issues.

Vendor partnerships represent a unique opportunity for the academic community to influence the design and availability of digital video services in higher education. An RFI was issued in October 1998 to identify vendors for partnership opportunities. In 1999, ViDe will be expanding its membership in 1999 and selecting vendors to collaborate in further digital video development. Video on demand development areas of interest to ViDe for the next two years include:

1. Robust video and audio file sharing across multiple vendor client/server platforms.

2. Storage, file management and playback for H.323 videoconferencing sessions.

3. Expanded metadata and visual object indexing, particularly cross-platform search and retrieval of digital objects.

4. Implementation and testing of MPEG-4

5. Development of distance learning applications and projects with digital video.

6. Incorporation of robust access features for disabled users, including closed and open captioning, signing windows and audio to text transcoding/indexing

ViDe invites your suggestions and comments on any development area. We'd also like to hear what you are doing with digital video today and what you want to do tomorrow!

Email your comments to: gr[illegible]rary.gatech.edu

Date of Publication: April 2[illegible]