Digital Video Servers: Storage Technology and Applications

Richard F. Annibaldi
Pioneer New Media Technologies, Inc.

Abstract

How can cable operators help prepare themselves for the introduction of compressed digital video into their cable systems? This paper describes the technologies available for one important area, mass storage used in digital video servers. With this information, operators can better evaluate their alternatives in relation to their planned applications.

INTRODUCTION

As cable system operators strive to implement the National Information Infrastructure in the coming years, they will face many new challenges. In particular, the use of digital compressed video requires that cable systems implement unfamiliar technologies. Although some compressed video sources will probably be provided by satellite, the opportunity exists to provide other sources either directly in the cable system, or within a multiple system interconnect. These locally provided sources of digital video will take the form of a digital video server.

The basic function of such a server is to provide appropriate streams of digital information to the subscriber of the cable system. Depending upon the specific services that the operator wishes to offer, each of these streams may be interactively controlled by a subscriber.

VIDEO SERVER ARCHITECTURE

Let's begin by examining a typical structure of a server. Although the exact form used may vary for each cable system, digital video servers will likely contain several common building blocks. The three most basic blocks are a central processing unit (CPU), an input/output (I/O) system, and digital storage.

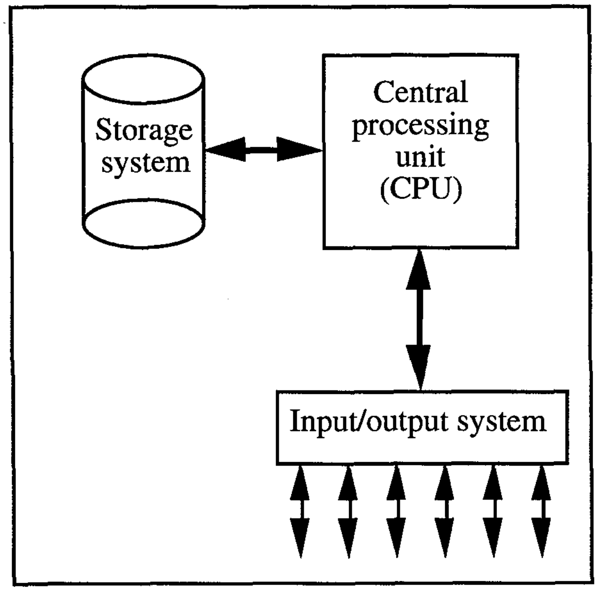

Figure 1: Basic Server Architecture

Our interest is in this third block used to hold large amounts of digital information. To allow for rapid movement of data into and out of the storage system, most devices use the Small Computer System Interface (SCSI). SCSI, pronounced "scuzzi", and the improved version, called SCSI-2, allow transfer of data at up to 20 mega-bytes per second (MBps), equivalent to 160 mega-bits per second (Mbps). Although this is substantially higher than the bit rates needed to support full motion video, remember that this interface must supply enough data to support multiple subscribers simultaneously.

DIGITAL VIDEO FORMATS

To evaluate storage systems, it is helpful to review the variety of digital video formats in use today. The storage system should be flexible enough to handle any combination of the video formats that will likely appear on the network.

Formats to be considered include:

MPEG1	Compressed full motion video at about 1.5 Mbps.
MPEG2	Compressed full motion video using several different profiles and levels at up to 60 Mbps - typically uses 3 to 4 Mbps for NTSC and 10 Mbps for HDTV
H.261	Video telephony and conferencing at bit rates up to 1.92 Mbps
JPEG	Compressed still pictures
Others	Proprietary standards including General Instruments' DigiCipher

Each of these formats may find a specific niche application on cable. For example, while movies would require an MPEG or equivalent format, interactive home shopping might use JPEG with freeze frames. Each format was developed based upon the desired compression algorithms, resolution and data rate. As a result, the actual data format stored on the media differs.

In fact, there is even more variety. With three profiles available at several levels each, MPEG-2 alone has many variations. The data may need to match a particular transmission format such as Asynchronous Transfer Mode (ATM). Although the video server processor can convert formats, storing the video in the final transmission format can increase throughput. The storage technology should be flexible enough to handle all these possibilities.

As a rule of thumb for later comparisons, consider the typical amount of data needed for a two-hour movie. With a 3 Mbps data rate, this movie requires 2.7 gigabytes (GB) of storage. Higher resolution, such as HDTV, would need much larger storage.

STORAGE TECHNOLOGIES

We can classify storage technologies into four basic groups: random access memory (RAM), magnetic tape, hard disk drives, and optical disc drives. Table 1 summarizes characteristics of each type described below. Because these types often complement each other, it is most likely that servers will use some combination of them.

Random Access Memory

The oldest, fastest and most flexible method of storage is Random Access Memory (RAM). Unfortunately, RAM is also the most expensive for large storage size. Nevertheless, its advantages make it desirable to use in conjunction with the other types of storage. After loading a RAM buffer from any other type of storage, the processor can quickly access bursts of data destined for multiple users.

Magnetic Tape

There are several digital video recording formats, such as D1 and D2, in use today. However, these machines record data in rigid formats and provide digital outputs primarily for duplication and editing. For more flexible data formats, magnetic tapes for computer applications are a better choice. Virtually all of these computer tape formats can store digital video. However, these tape formats share several disadvantages in video server applications. The two most important of these are long access time when positioning the tape to a random point, and the reliability effects of tape wear from repeated playback.

The availability of both hard disc and optical storage makes extensive use of magnetic tape in server applications less desirable. Nevertheless, magnetic tape can still be valuable for creating inexpensive archive copies of digital material that is no longer in regular use.

Table 1: Comparison of Storage Technologies

STORAGE TECHNOLOGY:	RAM	Magnetic Tape - Computer formats	Hard Disk	CD-ROM	DLD - Replicated	DLD - WORM
CHARACTERISTICS:						
Write capability	√	√	√			√
Erase capability	√	√	√			
Removable media for archive		√		√	√	√
Simultaneous outputs	1	1	1	1	4	4
Estimated total equipment cost per mega-byte per output (includes complete electronics, drive and all removable media, if any)	$25-$40	$20 on cartridge More for faster formats	$1-$2 on SLED $3-$16 on RAID	$50 for 1 to 6 disc changer Less for larger changers	< $1 on single disc drive <$0.10 on changer	< $1 on single disc drive
Average access time to any data location	< 1 µsec (processor access time)	50 secs for cartridge (tens of seconds for faster formats)	15.6 msec	< 500 ms within single disc	2 seconds within single disc	2 seconds within single disc
Sustained data rate for read operation in mega-bits per second (Mbps)	Limited only by processor speed	0.5 Mbps cartridge; up to 120 Mbps for others	16 to 40 Mbps; higher for RAID	1.2 Mbps	Up to 15 Mbps	Up to 15 Mbps
Typical storage increment in giga-bytes (GB)	Variable	0.13 GB per cartridge	Up to 1.6 GB per drive	0.54 GB per 4.75 inch disc	5.4 GB per 12 inch disc	5.4 GB per 12 inch disc
ADVANTAGES:	•Fastest	•Inexpensive media •Good for creating archive copies	•RAID improves access time and data rate	•Well suited for multiple copies	•Largest capacity •Well suited for multiple copies	•Largest capacity
DISADVANTAGES:	•Most expensive	•Very slow access time for random locations on tape	•Need RAID or backups for data protection	•Set-up cost for single disc copy	•Set-up cost for single disc copy	•Moderate cost for single disc copy

Hard Drives

This is now the dominant technology in the computer industry for large capacity storage, especially on personal computers. Hard drives combine the ability to read and write data at relatively high speed with non-volatility.

There are two general types of hard drive systems: the traditional Single Large Expensive Disk (SLED) and the newer Redundant Array of Inexpensive Disks (RAID). A SLED can now have capacities of 1.6 GB of data, with larger drives planned.

Groups of drives are combined with appropriate control software to form RAIDs. The different types of RAID configurations are classified as six different levels numbered 0 to 5. These "levels" do not indicate their relative merits; they simply identify different configurations that have different advantages and disadvantages. Table 2 summarizes the different levels. The characteristics of the various levels are improved speed and, more importantly, error detection and correction.

Table 2: Redundant Arrays of Inexpensive Disks (RAID)

RAID LEVEL:	RAID 0	RAID 1	RAID 2	RAID 3	RAID 4	RAID 5
CHARACTERISTICS:						
Data on original disk duplicated or mirrored on second disk		√				
Data striped across multiple disks using one byte per drive accessed	√			√		√
Data striped across multiple disks using full sectors on each drive accessed					√	
Error detection and correction codes stored on a separate check disk			√	√	√	
Parity interleaved with data and striped across several disks						√
ADVANTAGES:	•Increased speed	•Full redundancy of data	•Large data block efficiency	•Increased speed •Large data block efficiency	•Increased efficiency for small data blocks	•Allows multiple simultaneous writes
DISADVANTAGES:	•No error detection / correction	•Only 50% of disk capacity usable	•Unnecessarily redundant error detection / correction	•High overhead working with small amounts of data	•Slow writing of data due to shared check disk	•Most complex controller required

RAID redundancy for error detection and correction permits continued operation even after one drive fails. In addition, with many RAID controllers, an operator can actually replace a single failed drive without shutting down the system.

Besides the obvious approach of time sharing by multiple users, the operator also can configure some RAID levels to support simultaneous accesses to different drives. This would require careful handling of contention among users for the same individual drive. The specific advantages and disadvantages must be weighed against the desired applications to determine the best fit.

Optical Disc Drives

Compact Disc Read Only Memories (CD-ROM) and Digital Laser Discs (DLD) are the two primary types of optical disc technology. While CD-ROMs have been in use for many years, DLD is still under development. Details on DLD given in Table 1 are target specifications.

CD-ROMs: As the name states, CD-ROMs cannot be erased or rewritten. Like existing analog laser discs, CD-ROMs are replicated using a stamping process. Each 12 cm (approximately 4.75 inch) diameter CD-ROM can contain up to 540 mega-bytes of data. Drives are available for either a single CD-ROM or a magazine of six CD-ROMs. Changers for much larger quantities are also anticipated in the future.

The 540 mega-byte capacity of the CD-ROM is much less than the 2.7 GB required by our rule of thumb movie. Also, the CD-ROM data rate is only 1.2 Mbps. For these reasons, vendors are introducing several variations on CD-ROMs specifically targeted towards the video storage market. Some are spinning the disc faster to achieve two, three or four times the standard data rate. Others are developing products that differ more significantly.

As an example, Pioneer's α (alpha) Vision System uses high density recording and replicating techniques to place 2.12 GB on the same CD-ROM sized disc. The data transfer rate for α Vision is 4.7 Mbps, nearly four times the usual CD-ROM rate. The result is a disc which can store sixty minutes of video equivalent in quality to analog laser discs. In addition to the video, this disc stores two full stereo, or four monaural audio tracks. Other data can be included at a transfer rate of 130 kbps.

DLDs: A 30 cm (approximately 12 inch) diameter Digital Laser Disc can contain up to 5.4 GB of data, ten times that of a CD-ROM. DLD will support a variety of video formats including MPEG-2. A SCSI-2 interface carries the 15 Mbps data transfer rate.

A stamping process produces normal analog laserdiscs, called replicated discs. Currently planned playback drives for DLD will handle both replicated and Write Once Read Many (WORM) formats. WORM allows for digital storage of cable system specific content, while retaining all of the playback advantages of replicated laserdiscs. These advantages include non-contact reading to eliminate media wear, quick access to any location on the disc, and the ability to place seldom-used discs directly into archive storage.

Traditional analog laserdiscs contain video in either of two formats: constant linear velocity (CLV) and constant angular velocity (CAV). CLV allows for twice the storage capacity of CAV by varying the rotation speed depending upon read location.

By using the CAV format, however, DLD can use more than one pickup at a time. For example, locating a pickup at each of the four major compass points around the disc allows four simultaneous, yet independent, outputs. Figure 2 illustrates this four pickup drive, showing the four pick-ups or heads, each reading a different part of the data on

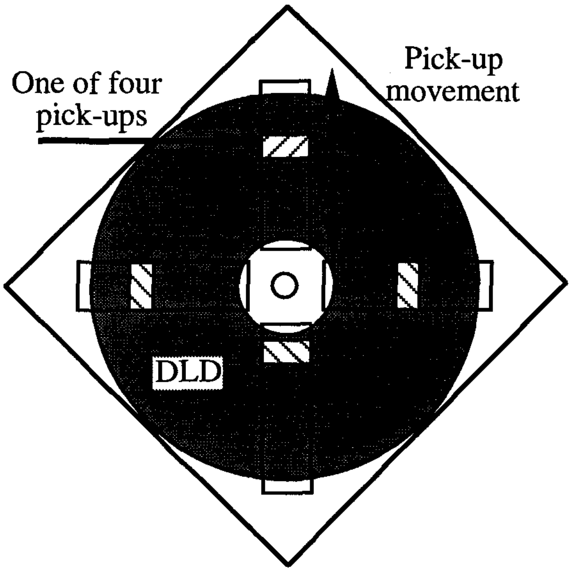

Figure 2: Four Pick-up DLD Player

the disc. Each pick-up moves across the spinning disc to the desired start position, from which it tracks the spirally recorded data.

In addition, there are also plans for automatic disc changers with two independent players and up to 252 DLDs. These changers allow easy access to huge libraries of information.

A variety of support equipment is also under development for DLD. First, a four channel MPEG2 video decoder allows use of the DLD drive as an analog video source. This makes for an easy introduction of the DLD into today's cable system. Once digital transmission begins, the operator simply removes this decoder from the headend.

Second, a four channel MPEG2 data synchronizer/multiplexer provides the steady data stream needed for reliable, flicker free video decompression at the set top decoder. Next, a combined encoding and authoring station handles preparation of both real time and non-real time data for recording. Finally, the WORM recording system writes the data onto the twelve inch disc.

CABLE SERVICES

Now that we understand the technologies available, we can apply this information to the cable system's video server. What services does the cable operator plan to offer? Let's look at each of these and see which technology fits best.

Near Video on Demand (NVOD) consists of providing the same movie on multiple channels with staggered start times. For example, one two-hour long movie could start at 8:00 PM on the first channel, 8:30 PM on the second, 9:00 PM on the third and 9:30 PM on the last. NVOD is ideal for first run movies because it can handle very large numbers of purchases using only a few channels. From our list of technologies, the four pickup DLD stands out as the best match for NVOD.

Video on Demand (VOD) provides individual access to a program. For example, each subscriber can purchase a particular movie, starting and pausing that movie whenever they wish. VOD requires rapid response to subscriber requests. Also, the number of programming choices to be offered to the subscriber has a large impact on the selection of technology.

Hard drives are a possible approach to VOD, although relatively expensive when offering many program choices. A DLD changer with as many as 252 discs can provide maximum choice, containing over 1300 giga-bytes of data. Because the DLD changer contains two players with four heads each, it can handle multiple accesses to this library. In addition, data retrieved from the DLD changer could be buffered in RAM while a different disc is loaded into one of the players for use by other subscribers.

High Definition Television (HDTV) can easily replace standard compressed video on the cable system. Both NVOD and VOD can work with these HDTV signals. The only requirements for the server are the higher bit rate and the specific digital format used for HDTV.

Interactivity can be very different from the full length feature films mentioned above. Films are stored sequentially in memory and require only occasional VCR-like controls to interrupt the normal flow. On the other hand, interactivity can involve random access to very short sequences, such as still frames of video or even short blocks of text. When interactivity requires many random accesses, hard drives are the logical choice. If video sequences are lengthy, however, DLD may have some advantages. In either case, the best approach is to store subscriber inputs in RAM to avoid disturbing the source data on hard disk or DLD.

EVALUATION STEPS

The following steps can help to determine the storage needs for a particular cable system. Simply answer each of these questions for your system:

1) Which services will you carry in the short term? What about the long term plans? Use these answers to match appropriate storage technology to each planned service. Most likely, you will want to phase in the new services gradually.

2) What is the expected popularity of each service? Use this answer to help identify the relative quantities of each storage type needed.

3) Which technologies can support the immediate applications without becoming obsolete when new services are later introduced? This will ensure that your system will grow and adapt to your needs.

Provide these answers to potential video server vendors so they can propose suitable solutions to your requirements. With your guidance and feedback, they can provide a much better system for you.

CONCLUSIONS

To be ready to face the challenges of the digital age, cable system engineers must gain an understanding of many applicable computer technologies. One such technology is the storage of large quantities of digital video and other information. The ideal storage technologies for any cable system will depend upon the services offered. Knowing your applications and the capabilities of various technologies should help in equipment selection.

ACKNOWLEDGEMENTS

The author wishes to acknowledge the assistance of Koh Uchiyama of Pioneer Electronic Corporation in providing source materials for this paper. In addition, thanks to many individuals at Pioneer New Media Technologies for helping to review this work.

REFERENCES

(1) Alford, Roger C. "Disk arrays explained", *BYTE*, October, 1992, pp. 259-266.

(2) Casey, M. "Disk arrays: Hype vs. reality", *Computerworld*, August 19, 1991, pp. 54-55.

(3) Ranade, Sanjay. *Jukebox and Robotic Libraries for Computer Mass Storage*. Westport, CT: Meckler Corporation, 1992

(4) Stone, M. David. "Hefty Storage on a Budget", *PC Magazine*, October 12, 1993, pp. 247-264.