The Application of Disk-Based Video Servers in News and Sports

By Jon Shaw

The pressures on broadcasters to cope with a rocketing demand for news and sports programs from a modern day audience intolerant of lateness, mistakes, and amateur presentation are immense; however, there is an increasing expectation that the new disk technology will dramatically ease the burden. This paper looks at the impact that available disk technology is currently making on the production processes of TV news and sports. The choices available are measured against the imposing demands of the broadcast environment, and the ability to provide an integrated solution is assessed.

News and sports are undeniably two of the most demanding areas within broadcast television.

With the falling price of disks, increased capacities, and the ability to group disks together in redundant array of independent disks (RAID) configurations, changing the recording medium from tape to disk is seen as the key to speeding up access and transfer, allowing for further automation and improving reliability.

Adopting disk-based systems represents a large step. They are seen as the way of the future, but what do they offer and why does their performance vary so much? Much of the controversy springs from the separate approaches of two industries—computers and television. The former sees servers for broadcast as an application of its existing file server hardware, whereas the latter may look for everything that can be done with tape, except better. Both approaches have their merits, but only through further development of RAID storage is there the potential to design video servers purposely built for the broadcaster.

The News and Sports Application

The demands placed on the system news and sports have similar applications. The worst case, where everyone wants to do everything at the last moment, happens all the time. The system must be engineered to cope accordingly. online disk-based editing and video servers bring new possibilities for streamlining, editing, and getting finished stories to air fast. But if the new technology cannot support the increased efficiency and speed demanded, there is little point in moving forward with it.

So Why Disk?

It has long been obvious that the process of getting finished stories to air fast would benefit from the faster access, lower maintenance and running costs, higher reliability, etc., offered by disk technology. For a while, insufficient capacity and redundancy kept the technology from widespread use. With the advent of RAID, both of these concerns have been overcome.

The situation now is whether to 1) tread carefully by taking a step-by-step approach, that is, start with a few small units and gradually build up a system of distributed storage as confidence and experience increases; or 2) aim for a total solution from the onset, removing tape from the process as quickly as possible. These difficult decisions impact directly on a broadcaster's ability to operate swiftly and efficiently and govern his ability to compete. Ultimately the approach taken is governed by the type of technology chosen.

If the broadcast environment were less demanding, the choice for storage would be largely based on amount of storage and number of users. In the computer industry, file servers incorporating RAID arrays, designed to provide large central shared storage and access to files for a number of online users, are used. Generally, the files are of a few kilobytes and should be delivered within a second of request.

In the broadcast environment the demands are more severe:

- With the inevitable trend toward compression in acquisition and transmission, it is becoming increasingly important to maintain the integrity of the signal through the production process. If compression is used, its quality should be as high as possible, and if any kind of processing (e.g., effects, multilayering, etc.) is carried out, there should be no resulting degradation to the original image (i.e., processing should be carried out at full bandwidth).
- The amount of storage required is vast, running at 75 Gbytes/hr for noncompressed ITU-R 601 pictures. Even if compression is used, perhaps allowing this figure to be divided by 5, the resulting number is still enormous. Video servers of around 500 Gbytes (0.5 Tbytes) are not uncommon.
- The data have to be delivered to each user continuously at 21 Mbytes/sec for noncompressed live video.
- Random picture access must be lightning fast, i.e., within 1/25 of a second.

The key demand of disks is to provide maximum access time. Grouping disks into RAID configurations does not necessarily offer the required performance. Although storage and bandwidth are relatively straightforward to achieve by grouping arrays of modern high-capacity high-speed disks, access time presents the greatest technical

Presented at the International Broadcasting Convention, (paper no. 18), in Amsterdam, The Netherlands, September 12-16, 1997. Jon Shaw is with Quantel, Newbury, Berkshire, England RG14 2NE.

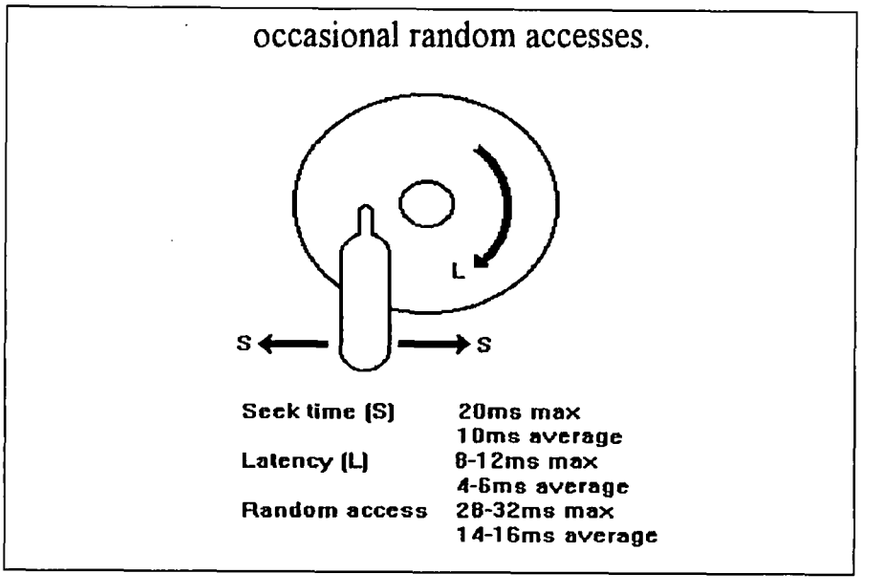

Figure 1. Access parameters of typical disk.

challenge and potentially the greatest benefits.

Disk Storage Choices for the Broadcast Environment

The whole nonlinear theme is based on the faster access disks offer over tape. Even so, mechanical constraints (an arm must swing to position the heads over the track, and the disk must spin to the start of data [Fig. 1]) dictate that individual drives cannot provide true random access—the ability to access any frame in any order at video rate. The time between frames is 1.6 msec—around 1/10 the average random access time of a fast disk. Typically, RAID stores can only continuously play clips if the frames are stored on adjacent tracks, making only occasional random accesses.

Whereas most video servers must run defragmentation and consolidation routines, those with true random access can run fragmented, continuously making random accesses. Defragmentation and consolidation routines may not be a problem for offline application but cause major problems for online application and a major concern for a video server. Its large store may take many hours to reorder the material; a potential crisis for a vital central resource.

For continuous, unimpeded access to video stored on disk, file size must be of the smallest useable size for TV; i.e., 1 frame. Thus every frame can be individually accessed in any order without the need to open larger files. Achieving this speed of access demands performance that is not available from off-the-shelf RAID arrays. It must be specially designed.

The vast majority of manufacturers involved in supplying disk products have either discounted the benefits or simply not attempted to build this type of store. This has resulted in a plethora of disk storage choices available to the broadcaster that offer certain benefits but are far from ideal. These range from the fully blown computer-type file servers, through to video disk stores using similar file storage techniques, but with modifications (SDI I/O, RS 422, etc.) to help slot them into the broadcast operation.

Some video disk stores are primarily aimed at replacing VTRs by offering greater speed and reliability but with little more in terms of functionality and their ability to offer a totally integrated solution. They are proving to be a popular choice because they can provide an adequate replacement for the VTR in specific isolated operations (time delay and caches in transmission being two of the most popular), and provide broadcasters with a relatively cheap way of trying out disk technology without making too great a commitment.

The limited functionality and performance of these video disk stores significantly hinders their ability to be used for different processes at the same time (e.g., editing and transmitting simultaneously). To expand any system based on this type of device, either to increase the amount of storage and users or to incorporate different processes, inevitably involves adding more units. The attractions of starting small with a relatively modest commitment melt away if later expansion brings increased complexity and compromised operation.

An integrated solution needs a storage device as large as is practical to avoid the need for additional units (avoiding the problems of copying, access, transfer, etc., associated with distributed storage). It requires flexibility to provide simultaneous functions and reliability and must operate comfortably within a multiserver environment. Such a device could be called a true broadcast video server.

The True Broadcast Video Server

An example of such a server is Clipbox from Quantel. This uses specially designed hardware to meet the necessary access requirements, discussed earlier. The basic building block is the Dylan disk store.

In this, an array of 20 disk drives is used as the storage for one noncompressed channel. The arrangement is for one drive to provide redundancy, so that full operation continues with no data lost should any single drive fail. The combined data rate of the remaining 19 drives amounts to over 95 Mbytes/sec, which is sufficient to allow for continuous random disk access while delivering a continuous 21 Mbytes/sec of data.

In Clipbox, up to eight of these stores are combined allowing their capacity to be shared between all video ports. Clipbox has up to 8 hr of noncompressed video storage. Also, compression (using Quantel's grid compression scheme) at 5:1, 10:1, or 20:1 is used to give up to 160 hr storage. Uniquely the store can hold a mixture of compressed and noncompressed material and supports up to 8 ports, or users, noncompressed or 14 with compression.

Every port on the server has the

ability to access any stored picture, in any order, at video rate (21 Mbytes/sec) regardless of the demands of other ports. This way the operation can continue regardless of the extent to which the store is fragmented or the order required for recording or replaying pictures—there is never any need to consolidate, de-fragment, or make copies. Each port is independent and can be used as a separate entity, allowing diverse activities to happen on one server.

The fact that pictures of any clip may be scattered, in frame increments, throughout the store is excellent for flexibility but would present problems for any external system needing to access material. For this reason the internal management database addresses every frame to present a group of them as a clip, or sequence of clips. Thus external addressing is resolved to the relatively simple matter of recalling a few large items rather than thousands of individual frames.

This technology is leading edge, resulting in a working product that offers the performance, flexibility, and reliability that can only be offered by a server whose ports can each deliver true random access.

Applying the Broadcast Video Server to News and Sport

Figure 2 outlines the various activities and processes of a typical tape-based news (or sports) operation.

The main activities are summarized as recording, distribution and viewing, editing, and transmission. By having a central store (Fig. 3), these areas are integrated, avoiding the bottlenecks associated with distributed hardware.

Recording, Distribution, and Viewing

In many cases, most material is recorded from lines, but it may also be supplied on tape or disks, or accessed from archives. The number of ports assigned to this task depends on the load. The moment that material is stored in the server, it is available to all ports and all online users instantly without copying.

Most disk-based systems suffer from fragmentation thus requiring a de-fragmentation or consolidation program be run on a regular basis to reorganize fragmented free space into usuable chunks. Failure to run these routines can result in failure to record an incoming feed at the critical moment even though the system has reported to the user that it has enough free space to do the job. Running housekeeping routines such as these involves taking the system off line, something which may be considered unacceptable in a news environment. On the other hand, a video disk store with true random access avoids these problems allowing recording to take place no matter how fragmented the disk has become.

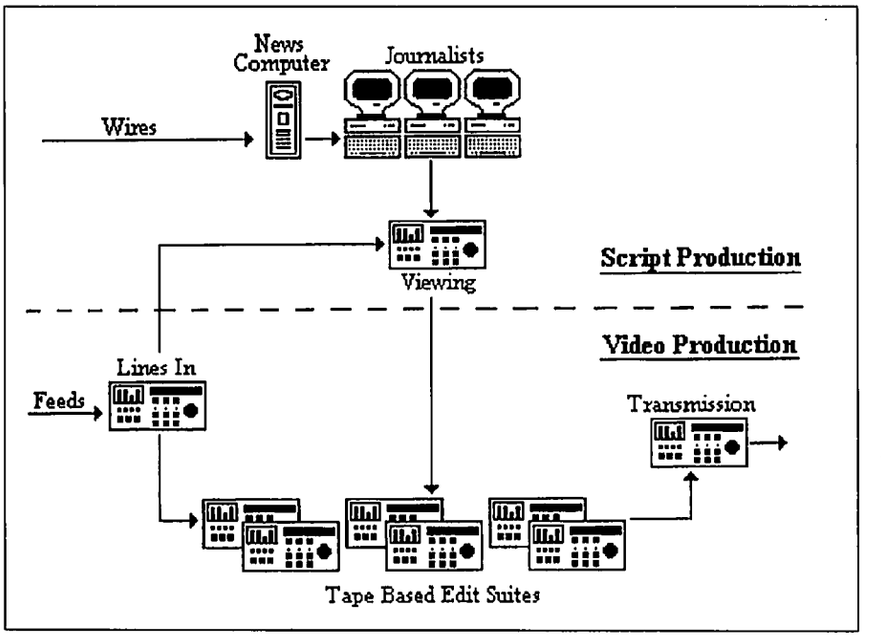

Figure 2. Traditional news/sport production process.

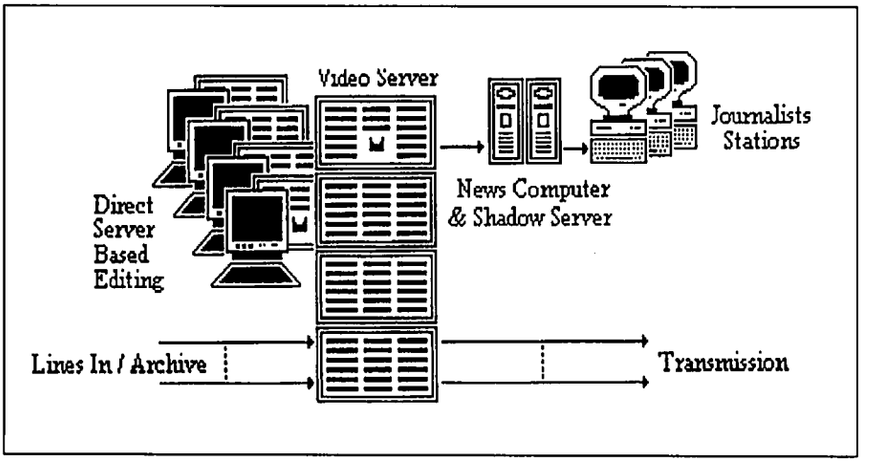

Figure 3. The news/sports production process based around a central store.

Editing can be started even while recording continues. Useful shots can be selected and instantly assembled into a rough edit, saving time in the final editing process. When operating close to sports and news deadlines, such flexibility is a real benefit.

Editing

To many, news integration must include an opportunity for journalists to edit video on their terminals. This can be accommodated by a shadow server, which is part of the newsroom computer. It mirrors, in highly compressed form, the material in the cen-

tral server, to make it available on the journalists' screens alongside text.

Clipbox can keep the shadow server in-step with the central server by a remote control protocol that identifies and allows the transfer of all, or selected, new material. This can operate from just one port on the store, but for heavy operation other schemes, such as direct recording in parallel with the main server, can be implemented. Journalists can then select the in and out points of wanted shots and assemble them into a running order. In other words, edit the video. Provision of audio is also possible. The decisions are returned to the Clipbox, where they are instantly conformed and available for replay with broadcast quality pictures and sound.

Despite the contributions from journalists, it is likely that "craft" editing will still be required. Quantel has pioneered online, nonlinear editing, developing an editor (Editbox) for news and post applications that depends on its own dedicated video disk stores for standalone operation. However, a single port of a Clipbox server can replace that store so that editing is done directly on its shared video.

A number of "diskless" Editbox edit systems can connect to Clipbox to provide a full range of online editing facilities. Unlike operating with tape, where cuts have been the only practical picture transition for news, it is just as easy to include more variety such as wipes, dissolves, and digital video effect moves. Here, voice-overs and graphics can also be added. Because the material is always stored in the single Clipbox, the result is always immediately available for transmission.

All cut edits exist as only decisions for the order of replay of the original material. For wipes, dissolves, and other multisource transitions, only the transition frames are recorded, adding very little to the stored material.

Transmission

Traditionally the replay of stories to air has been assigned to a separate, dedicated area with its own equipment. This gives a degree of protection to operations but also demands that finished stories be physically transferred to the transmission equipment. Now there is the opportunity of transmitting from the same server as the stories were edited and the original material was recorded.

What makes transmission from the same Clipbox technically possible is the independence of the ports. Operations on other ports cannot interfere, so transmission can proceed without interruption even though editing and recording may still be taking place on the same system. Not only does this save the cost of another server, but there are real benefits in operation. First the potential bottleneck in moving finished material to the transmission server is eliminated. This also removes the administrative load on the news computer to control and track the material. As soon as a story is edited, it is instantly available for transmission. This can be very important as late changes occur; stories are reedited and running orders changed.

The Broadcast Video Server in Practice

Our experience is that for most news and sports applications, a single Clipbox unit providing a single central store of video and audio for a number of simultaneous users is sufficient. (70% of those units installed are fulfilling this central role.) There are, however, occasions when several units provide additional redundancy or more storage and users. These servers must maintain local operations (e.g., post-production, transmission, etc.) and also interface to outside servers.

Because the ports are independent, the operations on interconnecting ports do not disturb other server-based work. Routines available in the remote control protocol also help; for example, if the servers are intended to hold identical information, then flagging changes in one enables them to be reflected in the others. Software to automate this process is relatively simple.

The implementation and use of the Clipbox has been extremely diverse, no two installations have been, or are likely to be, the same. There are a myriad of possibilities requiring flexibility. The common denominator in all these installations is the core hardware of the server.

Conclusion

There are many different ways for a broadcaster to adopt disk technology and use it within news and sports production; there are many products available that can offer improvements to specific areas of the operation; however, there are very few products capable of bringing the whole system together to completely change the way that these production processes are carried out.

The true broadcast video server, designed from the ground up for the broadcast environment, is the way forward and is already causing a revolution in the way that news and sport production is being carried out. There are many broadcasters enjoying the kind of benefits to their operation that have been described, and there are a number of broadcasters who have adopted the technology to use in, and alongside, existing installations. They know that in time as reliance on, and confidence in, disk increases they have invested in technology that will grow with them and ultimately allow them to completely revolutionize the way they produce news and sports.