

A Server-Based Post-Production System

By Bob Pank

The use of disks in video is spreading rapidly beyond individual machines and into servers. Several distinct application areas can be identified, each with their particular demands for bandwidth and access, which constitute major differences in performance — that of post-production being the most demanding. The design fundamentals for this server and configuration of the whole post-production system are reviewed, along with its potential benefits in a practical system.

Real-time digital video recording on disks started about ten years ago. Then small capacities and high costs limited the applications to specialist editing and graphics compositing. Tumbling disk costs, coupled with rapidly increasing capacity, have encouraged the development of off-line editing with low-quality compressed pictures, as well as ever longer stores, for full-quality, on-line use. As the development of individual machines for graphics and editing continues, video servers have become the latest area for disk applications. The requirements of post-production make it the most demanding of all application areas for servers, but the benefits for operation and in running costs have great appeal and therefore servers are set to shape the future of post-production.

This paper categorizes the general applications of video servers and then concentrates on post-production, using Quantel's Clipbox server and Editbox edit suite as a practical example of a server-based system.

Server Types

Although the technology is still in its infancy, there is already a range of servers available. A brief examination of these highlights the special requirements for a post-production server against those suited to other applications.

Server applications can be divided into four areas: video on demand (VOD), transmission, news, and post-production. Despite these all having large storage capacity and multiple channels, they differ widely in performance, as they are tuned to meet specific requirements. Recognizing and matching these is essential as the server represents a large section of technical equipment and can have a profound influence on the performance of the whole facility.

Four principal variables can be identified. The first two, the number of channels and the quality of pictures and sound, together define the total bandwidth. The second two, access speed and play time (how long any item runs before another random access is needed), are both a part of the overall access specification.

Storage requirements will also vary, but in most cases are defined by individual needs and budgets, as all types of servers can be now be built for many hours of storage. Reliability and the degree to which the system is protected against any failures will also vary according to customer needs.

Video on Demand

VOD is required to provide a very large number of channels, each of low or VHS quality, which together may demand a high bandwidth. Access speed can be slow, as any chosen item will be expected to start replay no sooner than within the next few seconds. In some cases this may stretch towards a minute or longer, in which case the access tends to be called near video on demand (NVOD). Once

accessed, play time will be anything from a few minutes to a few hours, so that overall access demands in the VOD server are very low.

Transmission

Typically, a smaller number of channels is needed, say up to 15, but these have to be presented with a picture and sound quality to match the actual delivery system. The NTSC, PAL, and SECAM color coding systems in use today are all forms of analog compression, and digital transmissions based on MPEG-2 compression are now rapidly starting up. The point is that some form of compression in the server is acceptable, provided it does not degrade the pictures and sound received after transmission.

A modest number of channels and "medium" compressed video together mean that the bandwidth requirements are relatively low. Access to items needs to be guaranteed to frame accuracy or at least to within a frame or two, as the system will often be controlled by a scheduling computer. Manual operation is quite different from being cued and given notice in advance by an automation system, so instant random access to any item needs to be checked. Play time can range from 5 seconds for a short spot to a few hours, so access demands are only moderate.

News

This application can represent a combination of post-production and transmission. The maximum number of channels will again be around the 15 mark, with analog Betacam quality generally being looked for; some compression is acceptable, making the total bandwidth only moderate. For editing, access requirements must be on a frame-by-frame basis with the play time typically down to the length of the shortest shot — a second or so, making quite high demands for access.

Depending on the performance of the server and the independence of its

Presented at the 137th SMPTE Technical Conference in New Orleans (paper no. 137-15) on September 7, 1995. An unedited version of this article appears in *Moving Images: Meeting the Challenges*, SMPTE, 1995. Bob Pank is with Quantel Ltd., Newbury, Berkshire, England RG 13 2 NE. Copyright © 1996 by the Society of Motion Picture and Television Engineers, Inc.

A SERVER-BASED POST-PRODUCTION SYSTEM

Table 1 — Server Performance for Applications

	Bandwidth	Access	Total
VOD	High (3)	Low (1)	4
Transmission	Low (1)	Moderate (2)	3
News	Moderate (2)	High (3)	5
Post-production	High (3)	Very high (4)	7

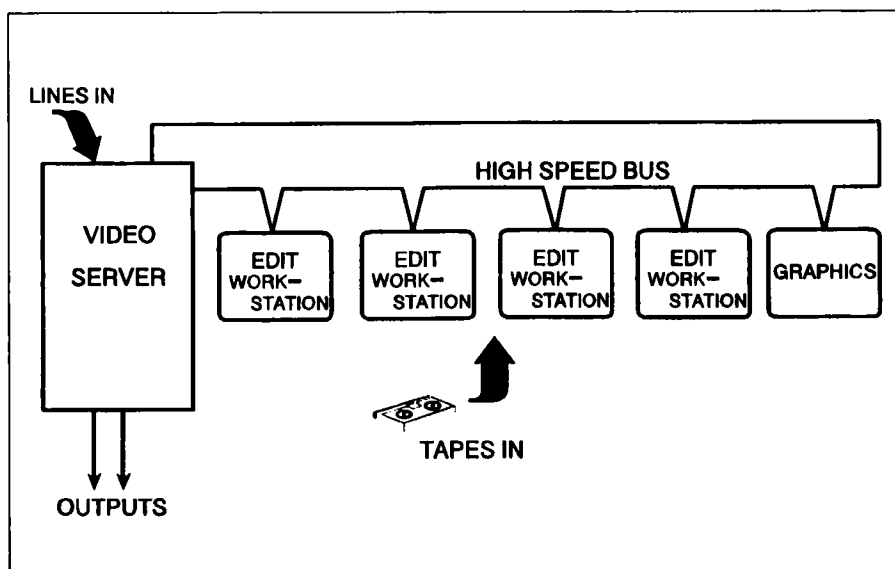


Figure 1. Video server network — computer ring configuration.

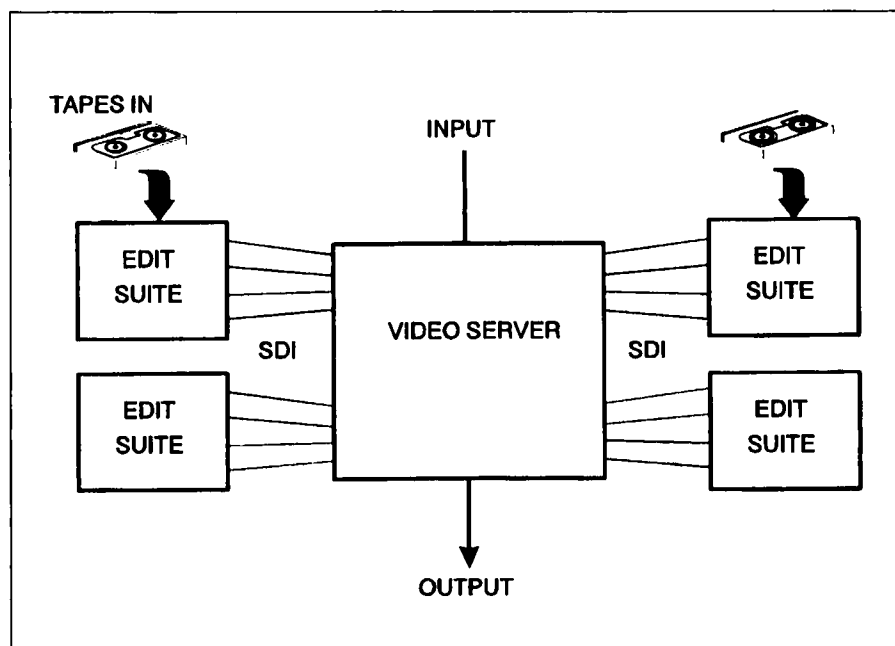


Figure 2. Video server network — star configuration.

channels it may be possible to use the same server to play out the stories directly to air. In cases where the uninterrupted operation of a port cannot be guaranteed, or the method of operation dictates, the finished stories would need to be copied to a separate device for playout.

Post-Production

The needs of the server are, in some ways, the opposite of VOD. In place of vast numbers of channels and VHS quality video, the post-production server requires only a modest number of channels, but each able to work with the full ITU-R 601 signal — preserving quality through to the edited master. This results in a need for high bandwidth. To cover the broad needs of editing the frames need to be played in any order, so access should ideally be completely random and at video rate (true random access) for each user.

The very high continuous bandwidth that has to be achieved, along with true random access to the store available at every port, creates a worst-case scenario, as there is a fundamental conflict in achieving high data rate and fully random access at the same time — something that is not a requirement of any of the other applications.

System Configuration

The details of how a server is used will always be a matter of individual requirements, but post-production is particularly demanding and so effectively limits the configurations that can actually work (Table 1).

Computer Ring

This is a copy of the ring networks commonly operated with computers (Fig. 1). While adding or removing machines on the ring provides an easy way to up or downsize the network, the configuration depends on the one ring highway to provide video for all suites. Considering the data rate and access requirements for editing multiple channels of ITU-R 601 pictures, no such highway is economically available today. Such configurations can be used when operating with highly compressed video, such as that used for off-line or journalist stations in news.

Star Configuration

Virtual VTR

Using an interface per channel, as in a star configuration (Fig. 2), means the transport of data between the server and the outstation "clients" is effectively shared and the data rate on any one interface is divided to a realistic level. One scheme is to use the server channels to substitute for VTRs in a conventional edit suite, a channel per virtual VTR. This allows traditional equipment to be used with disk stores but is not an efficient use of server channels. For example, four edit suites would require up to 16 server channels of full ITU-R 601 video — an impractical demand.

Integrated Edit Suites

The star configuration does allow the delivery of very high data rates from the server, but the needs of the traditional edit suites make excessive demands on the server's total bandwidth. The use of the channels as VTRs takes ideas from the tape era of post-production and applies them to the disk age. A more practical solution is to work with a suite originally developed for on-line disk operation (Fig. 3).

Editbox is an integrated on-line nonlinear suite that can be considered as comprising two parts. One includes video and audio input/output and processing as well as the editor's control interface, while the other is a true random access disk used to store the material to be edited and effectively supply all the sources and destinations for the editing process. The store provides a single continuous video channel but with the replay of its pictures ordered by the editing system. This way cut editing does not involve dubbing only transitions requiring multiple sources, such as dissolves and digital video effects (DVE), which are recorded back to disk. Clearly, the operation of the system depends on the performance of the disk store to provide one stream of video and audio.

Clipbox is a video server capable of simultaneously recording and replaying a number of channels of non-compressed video and audio on its disk store. As each channel has true ran-

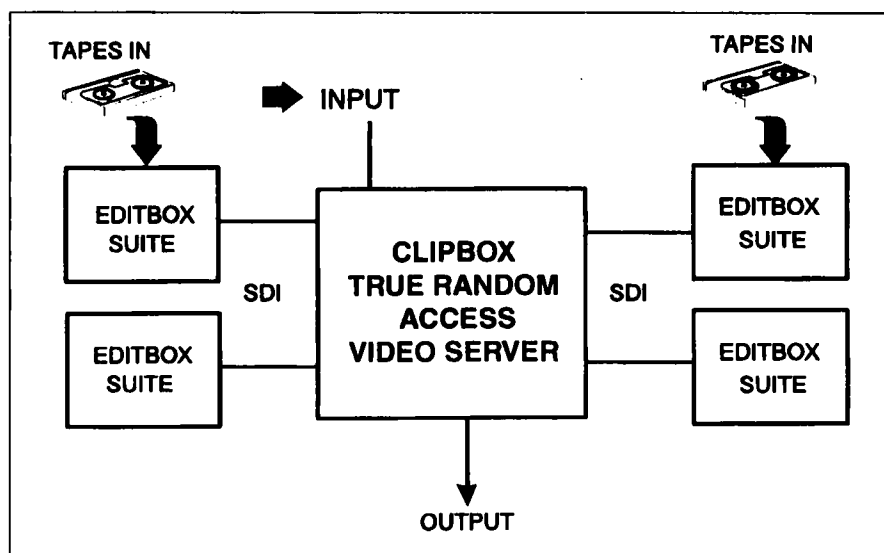


Figure 3. Clipbox server — editing and playout system.




Figure 4. Disk access time.

dom access to the contents of the store, it can substitute for the disk system normally dedicated to each Editbox, enabling it to operate directly with the contents of the server. This way one port serves a whole edit suite.

Design Fundamentals

True Random Access

It is relatively easy to produce a high continuous bandwidth while reading video from consecutive disk tracks in a linear fashion. On modern drives (Fig. 4) minimum track-to-track seek times are quoted as 2 msec, while maximum random seek times are

around 20 msec. Latency, or the time for the disk to spin to the required data, can add another 8 to 12 msec (for 7200 or 5400 rpm), meaning no useful data may be read for up to 28 or 32 msec. Taking the average as half the maximum, this is still well outside the TV frame interval of 1.3 msec and even compares with the whole frame time of 33 msec. For this reason disks have to operate linearly for a period, always moving to adjacent tracks, before building enough cache to allow time for a random access.

In such conditions fragmentation of data must be kept to a manageable level and must, from time to time, be

A SERVER-BASED POST-PRODUCTION SYSTEM

purged either by defragmentation, which reorders frames to allow them to play in a particular sequence, or by consolidation, clearing large continuous areas of disks for new recordings. Both take time proportional to the length of material, usually more than twice the running time, as both read and write operations are required to copy data on the disks.

Fragmentation is a random process and so introduces a degree of unpredictability into disk performance. With the larger stores of servers, the time to defragment or consolidate may increase to many hours — down time that may not be acceptable for on-line operations. As fragmentation cannot be avoided, this makes a strong case for building a server that can operate to full specification with a totally fragmented store. This is one of the reasons that the true random access, introduced with stores for editing systems Henry and Editbox, is included as a part of the Clipbox server specification.

Simultaneous Access

Perhaps the most obvious attraction of disk over tape-based operations is the instant access to all recorded material. In building a post-production server this degree of access should be extended to all ports, without restriction. This, coupled with the previously mentioned true random access, means that the same pictures and sound can be instantly and simultaneously output from any or all ports. Although the need to do this may not occur every day, the implications for efficiency are important.

Any material in the server can immediately be used by all connected suites without contention, so there is never any need to copy material to guarantee access, thus saving time as well as disk space. One application is work sharing, described in the next section.

Discretionary Compression

At the high end post-production facilities staff need to be able to edit without impairing the original signals — requiring the use of the full ITU-R 601 signal. Other productions, which are not as demanding, may favor the economies of compression. The server should be able to handle both forms of

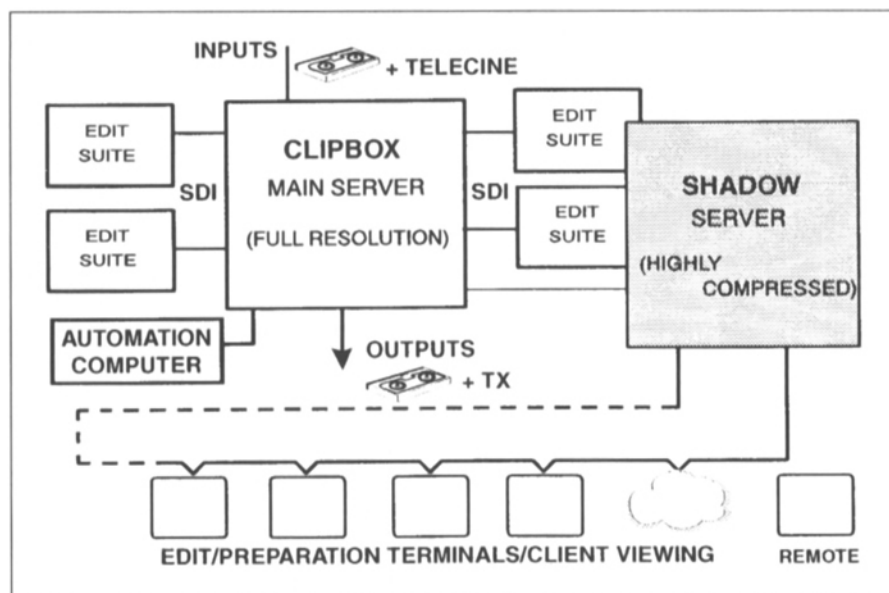


Figure 5. Post-production system.

video according to the needs of the production. The approach has been to build a server to operate with the full signal and then add the ability to handle compressed video — easier than upgrading a compressed system to handle noncompressed signals.

Security

Although disk drives are quoted as having mean time between failures (MTBFs) of over 1 million hours, some failures must still be expected. For continued operation under fault conditions and to guard against loss of data, an array of disks with redundant capacity is used. Data lost by a disk failure is regenerated from the error-correction system and stored on the new replacement drive. Statistically, the chances of two disks in the same array failing at the same time are so small as to be ignored. Further protection can be added by partial or total system duplication and is a matter of customer needs.

Management

Video stored in the server must be easy to find, use, and available to all users. The user has no interest in knowing the location of the stored frames but is only concerned with accessing whole clips or edited work. Internal management should allow whole clips to be searched by title, browsed, and viewed prior to selection.

Remote Control

A remote control protocol facilitates the direct use of other machines and systems with the server. True random access, simultaneous access, and the management system together simplify the protocol, making it easy for applications such as virtual VTR and play list automation to be implemented.

Connections

Post-production quality dictates the use of the full ITU-R 601 signal, but its bandwidth requirement is very demanding. The video alone requires an absolute minimum of 168 Mbits/sec, thus ruling out current offerings from the computer world. Asynchronous transfer mode (ATM) at 155 Mbits/sec clearly does not fit the bill, especially as, even at this bit rate, it can carry only around 70 Mbits/sec of useful data — network overheads and packetizing accounting for the difference. High-speed interfaces, such as high-speed parallel parallel interface (HIPPI), operate only over a very short distance and, using parallel transfer via multiway cable, would prove difficult to route.

The television industry already uses the serial digital interface (SDI) offering 270 Mbit/sec data rate over 200 m. A single coax cable, which is easily routable, can carry video, 16 channels of embedded audio, plus embedded

control. Besides using this as a connection between the server and the edit suites it provides the same connection as is already used on most modern TV equipment. As the server stands at the center of operations, being able to directly connect with other machines, such as VTRs and telecines, is an obvious advantage.

Benefits

The server-based post production system integrates editing and other facilities operating with a shared storage system. The advantages of this are described below.

Background Play-in and Play-out

The users of disk-based off-line editing suites will be aware of the time taken to copy recorded material onto disks prior to editing. By using the shared storage of a server, channels can be recording material for future jobs while suites continue editing other work. Likewise, completed edits residing in the server can be played out to tape or, possibly, direct to air. This makes use of remote control to operate the channels as a recorder or player, and takes advantage of their independence and free access to all stored material. Special occasions such as big sporting events can benefit from being edited while the event is happening, so when the final whistle blows, the edited highlights are ready.

Apportioning Storage

While some edits may require large storage, many others will not. There is a diversity of needs. For example, a four-hour server connected to four suites one day could offer three hours for a special event, with the remaining hour shared between the other three editing titles and commercials. The next day all four could each be used an hour for shows or serial dramas.

Work Sharing

Given that each suite can access all stored material means that several can be put to work on the same job simultaneously. For example, opening/closing titles and detailed effects work can be handled by the effects-oriented personnel, while the job of putting the rest of the show together is completed in another suite. There are no problems

of access, such as waiting for the other suite to finish with a tape, as the server allows all suites free access to all material.

Reduction or Total Elimination of Dubbing

Tape editing is all about dubbing, copying pictures and sound to make more tapes — B rolls, intermediate versions, rough cuts, the edited master, and distribution copies. The cost of tape, copying, storing, and maintaining the necessary VTRs, as well as the time taken, is enormous. With the post-production server it should be possible to go from unedited video to the replay of the final edit to air without any further copying of material.

A Practical System

The bandwidth and access requirements for a post-production server can be met and the server-based facility built (Fig. 5). In practice its success depends not only on the performance of the hard disk heart of the server, but also the hooks — controls, connections, and protocols — that bind the whole facility together.

The Editbox edit suites are connected via serial digital interface (SDI) using just two coax cables and operate directly with the video in the Clipbox server. They have no storage of their own. Other server ports are used for background load-in and playout using the remote control protocol driven by virtual VTR software run on a PC. Inputs from frame-by-frame devices, such as telecines and animation systems, can all be accepted through a single port and their results saved as individual clips in the server. If outputting directly to air the server's remote protocol can be driven by a scheduling computer to operate on a selected port.

Another facility allows a secondary "shadow" server to replicate the contents of the main store. Applications include the running of a 100% backup server or maintaining a compressed version of the main. The latter could be used to run a secondary system of terminals able to make off-line, rough-cut editing decisions (also useful for journalist terminals in news) or again to provide remote client viewing of editing progress via an integrated

services digital network (ISDN) link. Interestingly, the computer ring system configuration is appropriate here in off-line operation.

The success of the facility depends heavily on the contention-free and reliable continuous operation of the server. This goes hand-in-hand with remote controls and "hooks" to integrate an appropriate range of other equipment with the server to bring the full benefits of a server-based post-production system.

Conclusion

The current applications of video servers include video on demand, transmission, news, and post-production. The latter, with its requirements for true random access and noncompressed video, proves the most demanding to achieve from a disk store. Taking full advantage of the opportunities offered by servers involves the use of nonlinear, disk-based edit suites rather than VTR designs which, for speed and access, are directly connected via the TV industry's serial digital interface. This offers new methods of operation where all suites share common storage and edit directly on the store.

THE AUTHOR



Bob Pank is technical communications manager for Quantel Ltd., where he has enjoyed giving technical support for products and following the progress of digital television for over 20 years. Pank has contributed a number of papers to technical conferences and is the editor of *The Digital Fact Book*.
