Asset management in news and broadcast production environments

BY C. JASON MANCEBO

Since completing its digital conversion, the newsroom at Danish Broadcasting has enjoyed the benefits of integrated news production with a digital asset management and archive system based on SGI StudioCentral Library and SGI Media server for broadcast. Photos courtesy SGI.

Asset management is the latest buzzword. A lot of people are talking about it, and they all think they need it. But what exactly is asset management?

First, let's deal with a more basic question: What exactly is an asset? In news and broadcast production, an asset comprises two parts: content (or essence) and metadata. Content is the file (or set of files) that holds the digitized program material (for example, an MPEG-2 or DVCPro .dif file of a video clip). Metadata is information that describes the parameters of the content (bit rate, television standard, file format, etc.). In some cases, an asset may consist entirely of metadata with only pointers to other assets that hold the actual essence. An example of such an asset is an edit-decision list (EDL), which contains no actual program data but instead has pointers to the EDLs assets along with the EDLs in and out points and transition and overlay specifications.

In 1996, an asset management system was a computer program that consisted of a database with information about a physical tape. This information included the title or subject matter, the date shot or recorded, the photographer and/or producer, the physical location of the tape in the library, and perhaps (if the system was advanced), the barcode number. The term "metadata" was not a part of our daily language back then, but metadata existed and was being used. Typically, it consisted of handwritten notes of a shot log, or a computer printout of the EDL. But, in a mere five years, the industry has moved from being firmly rooted in magnetic-tape media to being engaged in the third generation of the digital transition: the tapeless environment.

Generations

The first stage of the transition started with the migration from magnetic-tape equipment to video servers and proprietary computer-based nonlinear editors. The first of these systems allowed more flexible editing as well as lower cost of ownership and operation. But the question is: Did they allow better management of digital assets?

Unfortunately, the typical equipment that led this transformation was little more than a tapeless VTR. In fact, most video servers or digital disk recorders (DDRs) of the era used storage that was similar to the videotape medium it was replacing. Videotape has a continuous area of magnetic particles on which to record data. There is no inherent structure to the medium; rather, the data provide the structure. Similarly, these first-generation devices used identical techniques when recording to their hard disks. They wrote data to a raw partition, an area of the hard disk that, like videotape, has only magnetic particles capable of recording the raw video and audio data. The raw partition does not have an inherent structure – a file system accessible to the computer's operating system. The data provide the structure. The downside to this is the lack of interchange capability.

An asset comprises two parts: content (or essence) and metadata.

Without a file system, there is no file. And with the raw partition defining the format of the video, the ability to transfer data to another system is severely limited. The most flexible of these systems were able to save the data to removable storage media. Still, the best that could be hoped for in an asset management system using this technology was the tracking of a physical asset: the videotape.

In the second generation of the transition, generic PC-based computers surfaced as the new hardware platform for much of the next-generation technology. Its intrafacility interchange capability provided an excellent gain in efficiency. Rather than using analog or SDI video as the medium of transfer, these second-generation devices are connected using standard Ethernet networks, albeit to homogeneous devices. Disseminating content via computer networks illustrates the concept

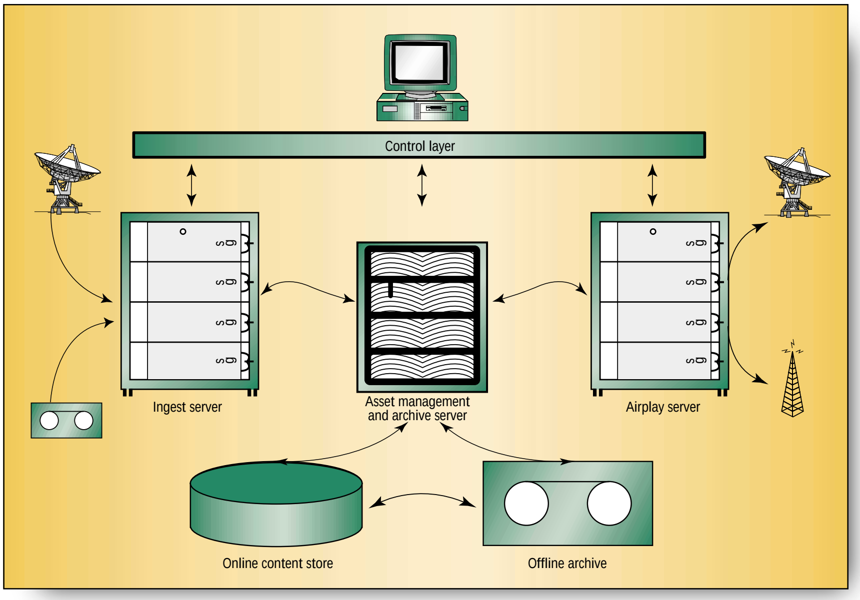

Figure 1. This block diagram shows the workflow of digital media asset archiving and retrieval used in broadcast applications.

of "distribute data, view video."

There are two associated actions—distribution and viewing—that you can do with video. Rather than keep these as a unified process, the "distribute data, view video" concept removes the bind between the two. Unless you need to view the video, you can distribute it over a network without viewing it and without using a time reference.

This model is superior for two distinct reasons. First, distributing the video as data does not require encoding or decoding, and therefore avoids the degradation of quality associated with these processes. Secondly, since you can distribute the video data without using real-time references, you can exploit the characteristics of high-speed data networks. It is now very common to have 100 Mb/sec, 1000 Mb/sec, or ATM OC-3 data networks within (and sometimes between) facilities. File-transfer times on these networks are several times faster than the video clip's total run time (TRT). For example, a 30-second file of 25Mb/sec MPEG-2 and its associated protocol overhead would transfer, via gigabit Ethernet, approximately 30 times faster than its TRT. By the same token, a news story with a TRT of one minute would transfer in approximately two seconds.

A StorageTek L700 and SGI 2100 server are key to the scalability of the digital asset archives at Danish Broadcasting.

Now, in the third generation of digital transformation, we move to use the "distribute data, view video" concept between heterogeneous devices. This means sharing assets not only between devices linked by high-speed local networks, but also on wide-area networks (WANs) between facilities of a station group, global WANs, or global public networks (the Internet).

There are two common misconceptions about network-based transfer and distribution. These misconceptions, and the factual explanations to dispel them, are listed below.

Misconception #1: The lower the bandwidth of the network, the lower the quality of the video.

The facts: The quality of the video is not related to the bandwidth of the network.

Since network-based distribution is file-based rather than real-time-stream-based, the quality of the video is determined when encoded or recorded. A clip recorded at 50Mb-per-second MPEG-2 will always have the same characteristics unless otherwise acted upon by further compression or transcoding. Simply transferring a file has no effect on the quality.

Misconception #2: The cost of a WAN connection with the bandwidth to transfer the file is prohibitive.

> ***The industry has moved to the third generation of the digital transition: the tapeless environment.***

The facts: The time required to transfer a file becomes a business decision. Measured leased lines can be an effective strategy for those who may have a sporadic rather than consistent need for faster-than-real-time file transfers. Examples may be a breaking news story or an immediate post-production session. At other times, when immediacy is not crucial (perhaps news stories for the next day or digital dailies), transfer can take place in slower-than-real time over a monthly lease line with a much lower bandwidth capacity and much lower cost. The key is that the network must support the requirement for data availability at the remote location rather than the compression format's requirement for real-time availability.

Digital asset management/ archiving

With the ability to transport many digital files over a LAN or WAN by many different workstations serving many different functions (such as editing, graphics, and acting as the ingest and airplay servers) comes the exacerbating task of managing the digital assets. Imagine taking a clip from the ingest server and sharing it with three NLEs as well as a graphics workstation. The original clip is now a contributor to five different clips. However, when the task of managing these assets is left to the asset management system, this once-daunting challenge at last becomes feasible (see Figure 1).

By using the asset management system at the center of the workflow as shown in Figure 1, all devices must check assets into the management system before other devices on the network can use the assets. High-speed networking should be used, and the speed of the network should be selected based on the format and compression (or lack thereof) used. A typical 30-second commercial spot using an MPEG-2 4:2:0 file at 4Mb per second (a typical playout format), would take approximately three seconds to transfer over a T-3 (45Mb-per-second) network between facilities, whereas an uncompressed 1920x1080i 4:4:4:4 file at ~250Mb per second would need nearly 25 minutes to transfer over the same network. Clearly, the speed of the network you use is a business decision as well as a product of the type of work in which your facility is engaged.

The speed of the network should be selected based on the format and compression (or lack thereof) used.

One additional item to note is the case of high-bandwidth, high-latency networks. When you use a network with an effective bandwidth greater than approximately 4Mb per second and a latency greater than approximately 15 msec (often referred to as long-fat-pipe networks or LFNs) in conjunction with an application that uses TCP/IP, such as FTP, you must be sure to select a device that supports RFC 1323. RFC 1323 is a TCP extension for high performance, which allows the TCP window size to scale. In LFNs, as the RFC reads, "TCP performance depends not upon the transfer rate itself, but rather upon the product of the transfer rate and the round-trip delay. The window-scale extension expands the definition of the TCP window to 32 bits and then uses a scale factor to carry this 32-bit value in the 16-bit window field of the TCP header (SEG.WND in RFC-793). The scale factor is carried in a new TCP option, window scale."

Without operating-system and application support for RFC-1323, transfer times over LFNs will be severely impacted, and the added capital outlay for the high-bandwidth network will be for naught.

Danish Broadcasting operator shown here performing ingest using Keyvia's (formerly Keops) Key-MediaWorks – creating and checking in multi-resolution assets to the asset management system from VTRs and satellite.

Asset sharing

Typically, several different digital-media data formats are used during the television production-and-broadcasting process. The number of formats will continue to increase as more video-compression schemes and file formats emerge and pervade the industry. Digital media will also continue to exist in several different media servers or file servers within a facility. Some content will be stored online, and other, less-frequently-used content will be stored in archives or in other types of near-line storage. This creates the need to search and access content, regardless of its type or physical location.

Searching can and should be extremely flexible. A system that supports data models is critical for a flexible, working system. A data model is a capability that allows the

Four SGI Media Servers for broadcast (purple rackmounted systems), and RAID-3 storage by Ciprico provide dual redundant systems for newsroom and airplay operations at Danish Broadcasting.

data structure of the asset to be defined. The most-flexible asset management systems will provide typical data models for common media-file formats. But they should also allow for user-definable data models. This ensures the systems will interchange with current and future file formats (see below).

Additionally, a critical factor that enables content sharing is a defined file format. Without one, there could be no interchange of assets between applications in a heterogeneous environment. While at least two key manufacturers attempted to urge the acceptance of their proprietary or wrapped proprietary file formats as open standards, the advanced- authoring-format (AAF) file [offered by the AAF Association (www.aafassociation.org)] and the MXF file format [offered by the Pro MPEG Forum (www.pro-mpeg.org)] are two excellent examples of industry working for the common goal of true file interchange in the most flexible and suitable format for the respective segments of the industry each organization represents. The AAF format is intended for editing and content-creation users and the MXF file format is aimed at streaming, ingest, and transmission uses. The goal is not only to exploit file interchange between heterogeneous devices by diverse manufacturers but to have AAF and MXF files interchange as well. This means that an editor could create an AAF format file using an NLE in New York and check it into the asset management system. Doing so would allow the transmission facilities (each in a different region of the country but connected via a network) to request the same file. However, they could each apply their own MXF filter to the original file. These filters enable the creation of a new file by selectively choosing and applying metadata within the original AAF file to the new MXF file. A typical use would be an AAF file that contained the metadata of the start time of a program to be played to air. This metadata would likely be in coordinated-universal-time format. With a user-defined AAF-to-MXF filter that applies the start time – -7 hours for central time or -9 hours for the user in the Pacific time zone – the metadata created in the new MXF file is customized. While this is an extremely simple example, one can see the potential of the AAF and MXF files and their filters. The key to all of this is the information quarterback: the asset management system.

Devices that do not have standard, native file formats limit flexibility and choice, and reduce efficiencies. Employing a standard information-technology infrastructure unlocks the key to a world of flexibility and lowers costs. Open-system file servers, the latest in high-speed networking, high-performance operating systems, and file systems are all examples of technologies employed by forward-thinking broadcasters managing and delivering their content as data. As such, they enjoy reduced capital outlay, ease of repair, and greater access to parts and service – all economies of scale.

Without compatible products using open and accepted standards, asset management by itself will do nothing more than allow you to manage your homogeneous islands of content. In this scenario, a user is able to query the metadata that is available on the local system or perhaps the local facility – a moderately interesting exercise that offers very little return on investment (ROI). The added value of asset management is the ability to share valuable assets, locally or globally.

Devices that do not have standard, native file formats limit flexibility and choice.

Archiving

As the transition to an asset-centric environment proceeds, the asset archive becomes increasingly important in the digital news environment.

The digital archive can store – on RAID arrays or on computer tape – all of the footage that has passed through the facility. The shooting ratio for most news/documentary productions is 50 to 1, even without reuse of final story footage, and the resulting archive material is an invaluable resource for the newsroom in the creation of future stories. The AMS performs a vital role in the management of the archive. The high-quality, high-resolution footage can be moved to less expensive, offline-storage formats such as data tapes or DVD-ROM. Metadata from the archived footage can be kept online in a database. Materials that are likely to be reused can be duplicated in low-resolution versions and kept available in online or near-line storage. This allows queries on the metadata and viewing of low-resolution versions of the footage via LAN, WAN, or public networks using standards-based, streaming-media technology.

The asset management system must support different modes of operation for its archives. Near-line and offline or archive storage can be provided by hierarchical storage- management (HSM) systems that provide seamless access to media contents by quickly bringing low-resolution footage onto a disk cache. One example is to use an HSM with a virtual file system to move infrequently accessed files to tape while keeping the most-often- used content files on a disk cache. In this case, streaming a low-resolution version of footage can initiate the HSM system to bring the footage from tape to disk. The high-resolution format of the footage could be archived automatically, or the direct approach of explicitly moving it to the offline archive manually could be used rather then letting the HSM decide when to move it. Furthermore, the facility administrator may want to control the specific tape or archive on which it is placed (for example, grouping all of the footage from a particular location together).

Making the transition to the digital environment seems a daunting task.

At first glance, making the transition to the digital environment seems a daunting task. With further exploration, one finds the current state of affairs an excellent indicator of the efforts manufacturers are making to provide truly open systems that will fulfill the promises of the digital transition.

While asset management systems encompass the entire workflow of a broadcast facility, not limited to acquisition and transmission servers as well as the database server, a flexible and effective asset management system includes high-bandwidth connection to content creation and editing seats, automation systems, and online- and offline-asset archives. This asset centric system depends on a strong API and software bus to unite the entire environment into a highly productive, well-connected and efficient workplace. Such a workplace saves time and money by accomplishing goals in four areas:

Content sharing and repurposing

- Decreasing the duplication of efforts to create or gather footage that the station or station group may already own
- Providing potential additional revenue streams by easier cataloging, tracking and versioning of assets and finished stories

Increased flexibility and creativity, allowing faster production

- Allowing access to all levels of personnel in local or remote facilities and increasing productivity and creativity

Decreasing capital costs of editing systems

- Allowing fast access to metadata and low-resolution versions of footage for the creation of rapid virtual clips and EDLs

Future-proofing capital investments

- Supporting data models and open file formats such as AAF and MXF
- Being flexible and scalable enough to work with existing technologies and future technologies that might be added to scale system capacity

Bearing in mind the above benefits, you must weigh the risks discussed and determine the ROI by carefully considering your expectations. But your thorough preparation will be rewarded with a system that satisfies both users and management. ■


From presenters to behind-the-scenes editors, all the members of a modern digital news department have desktop access to an asset management system on a daily basis.

C. Jason Mancebo is senior technology manager for the Media Industry division at SGI in Mountain View, California. He can be contacted at mancebo@sgi.com.