White Paper

A Data-Centric Approach to Cinema Mastering

Thomas J. True

1.0 Introduction

Advances in high-resolution real-time digital technology have revolutionized feature film production. From acquisition, through post-production to distribution, the cinema mastering workflow has become one based on the capture, storage and manipulation of high-resolution digital images. And, since these digital images represent the intellectual property and value of a production, it's of utmost importance to protect the integrity and quality of this data at every step in the production pipeline.

As filmmakers begin to use digital technologies throughout the production process, developers of hardware systems and software applications used by the Film industry are challenged with creating new products and methodologies that raise the bar in terms of the quality, and productivity. As such, there are significant technical requirements that must be met in order to support this workflow. These requirements include: (1) highest image quality, (2) high-bandwidth data transfer, (3) flexible SAN-attached storage, (4) security, and (5) archive capabilities. This paper presents a data-centric view of the production workflow and a summary of the latest technologies available to realize this workflow within a production facility. This data workflow is pictured in Figure 1.

2.0 Cinema Mastering Workflow

2.1 Acquisition

The first stage of the workflow is the acquisition or ingest of source material. Whether captured on film, or with the use of a digital camera, the role of the director of photography (DP) has extended further into the production process than ever before. Once the motion images are captured the goal is to create a digital intermediary using the highest quality images possible. This digital intermediary forms the basis for the production. These acquired images are stored as data in a format easily manipulated by the post-production tool suite. This universal mastering format offers the highest possible spatial resolution and a wide aspect ratio that permits all final distribution formats, as well as digital dailies and lower-resolution preview proxies, to be generated from a single copy of the original image data. This data enters the workflow from four principal sources:

- Film scanner/telecine
- Digital camera
- Image data
- Data recorder/VTR

Because the scope of this paper is cinema mastering, the focus will be primarily on data resulting from telecine transfers and digital HD cameras.

The telecine transfer process digitizes selected shots from the original negative or an interpositive. An EDL from an offline session dictates exactly which frames are scanned. This process generates DPX (SMPTE 268M) files, each of which contains a 10-bit LIN/LOG RGB image up to a 2K (2048 x 1556) resolution, along with a header and can be as large as 12MB in size. Once scanned, these image files are transferred to the digital archive, where they become immediately available for post-production work. Real-time transfer of these images requires approximately 300MB per second of sustained bandwidth. Meanwhile, the industry is moving toward the acquisition of 4K (4096x3038) film plates, with the goal of moving 4K data with 2K proxies through the entire workflow pipeline. This change quadruples the amount of data generated for each film frame.

In addition to the traditional telecine transfer of images from film, image data may also be acquired through the use of digital camera technology. This option has become more prevalent as digital film cameras that support the direct capture of 10-bit RGB 4:4:4 images have become available. In the past, HD video cameras were used, with the results stored as HDCAM™, D5 or D6 material for later transfer to the post-production system. This use of HDCAM and D5 has limited the data captured to a compressed YUV 4:2:2 data format. And D6, while uncompressed has still limited the captured data to YUV 4:2:2. New digital cameras, specifically designed for film work, now permit the direct capture of 10-bit 4:4:4 RGB, with the uncompressed results stored immediately on disk for later transfer to the digital archive. This technique requires the support of dual-link SMPTE 292M HD I/O and a sustained disk bandwidth of 300MB per second. In addition to the image data, time code information and other ancillary data must also be captured.

2.2 Post-production

The second stage in the workflow is post-production. Post-production is the series of processes that are performed on the original source material to produce the single final set of images from which all distribution formats are produced. In the digital production pipeline, software tools running on general and special-purpose computer systems and workstations implement these processes. Post-production processes include:

- Color correction/color grading
- Dust busting, grain matching and noise reduction
- Editing and compositing applications
- Paint and effects applications

During post-production, high-bandwidth, low-latency data communications are required between the digital archive and post-production systems, so that devices operating on the data can access it directly within the digital archive without making local copies. Making copies of the data introduces a risk to both the integrity and security of the data while at the same time creating multiple versions that subsequently need to be tracked.

Post-production software applications may indeed extract the image data from the larger DPX files and buffer it within an application for direct manipulation. Since during the process the data has been extracted and converted to another file format for easier manipulation, it is not a true copy. When an operation on the data is complete, it is imperative that the core image data is returned to the archive file, and each frame and the metadata associated with each image is updated to reflect and describe any changes that were made.

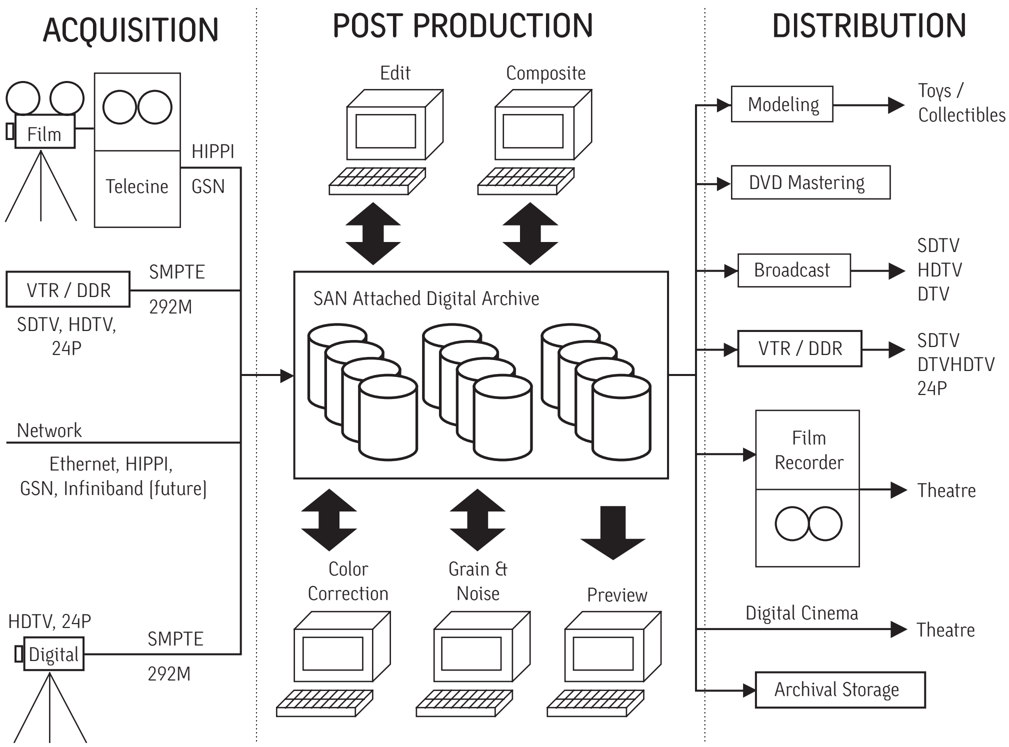

Fig. 1. Cinema Master Workflow

2.3 Distribution

Once all post-production processes are complete, the single digital master is ready for distribution. In today's environment, that distribution includes not only the final feature film recording but also:

- DVD
- TV broadcast – HD and SD
- Data archiving
- DTV
- Digital cinema

The data may also be repurposed for the creation of a sequel or for the creation of toys and other merchandise. These activities would utilize the same single resolution-independent universal digital assets. Finally, this single digital master, along with any ancillary data, is moved to nearline and then offline storage for later retrieval and reuse.

Since Tjabben and Bancroft originally proposed the digital production pipeline,[1] advances in technology have made the pipeline a reality. The remainder of this paper will outline the requirements of this digital pipeline and explain how current technologies can be employed to realize this pipeline within a production facility today.

3.0 High-Speed Interconnect

The transfer of 2K and 4K image data within a production facility with the goal of real-time or better transfer rates requires high-bandwidth communications at every step in the cinema mastering workflow. At 12MB per frame, 300MB per second is required to transport 2K data in real time. Moving to 4K data in the future will require more than 1GB per second of bandwidth for real-time transfer.

1. H. Tjabben and D.Bancroft, "Digital Film Mastering" SMPTE J., 859:864, December 1999.

3.1 HiPPI

One technology that has been effectively used to implement high-speed data transfer within the production workflow is the High Performance Parallel Interface (HiPPI) (ANSI HiPPI-PH X3.183-1991, HiPPI-FP X3.210-1992). HiPPI transfers 100MB per second using a 32-bit cable and 200MB per second using a 64-bit cable and can be used as a point-to-point connection between two systems or with a crosspoint switch in a LAN configuration. Using HiPPI, effective transfer rates for 2K data have been limited to 4 frames per second. At this rate, the transfer of 2K material between the telecine, storage archive, and post-production workstations has become a bottleneck within the workflow. Using a HiPPI connection between the telecine and data archive can take several days to digitize all the shots for a complete feature production. In the past this bottleneck has also prevented the full implementation of the digital workflow such that only shots that require special effects work are digitized, whereas the remaining shots of a film are edited using traditional film techniques. To improve the performance for telecine transfers, second-generation film scanning equipment implements Gigabyte System Network (GSN) (HIPPI-6400-PH, ANSI NCITS 323-1998) and High Speed Data Link (HSDL) to achieve transfer rates between 14 and 24fps for 2K data. GSN is the highest speed, highest bandwidth networking topology available today that is standards-based.

3.2 GSN

GSN (HIPPI-6400-PH, ANSI NCITS 323-1998) provides a full duplex point-to-point link capable of transferring up to 800MB per second over optical cables up to a distance of 75 m or over copper cables up to a distance of 40 m. Using TCP/IP and SCSI protocols, data is transferred error free in 32-byte flow-controlled micropackets. In the post-production workflow, GSN has demonstrated transfer rates approaching 24 frames per second. GSN interfaces can connect two computer systems directly or multiple systems in a switched network. GSN switches and bridges form data concentrators that permit GSN to be connected with Fibre Channel, HiPPI, SONET, Gigabit Ethernet, and other data communication technologies.

3.3 HSDL

Whereas HiPPI and GSN require additional infrastructure to implement high-speed data communications, HSDL[2] uses existing infrastructure within a post-production facility. HSDL, or High Speed Data Link, builds on the SMPTE 292M data model and SMPTE 374M 1080/4:4:4 dual-link serial digital video transport specification to move 2K images within a facility at up to 19fps. This alternative approach uses standard digital video interface connectors and as such can transparently use existing infrastructure within a facility. While HSDL is not a standard, it is being adopted by many manufacturers as a high-speed, point-to-point device interface.

HSDL moves data as contiguous frames encapsulated in a video format. Doing this eliminates the overhead of transferring individual images as files, each with header information that must be read and interpreted. It also simplifies buffer handling. And because HSDL is a video format, it can transparently pass through most HD devices, such as routers, being processed only on devices that support it.

HSDL encapsulates 2K image data within the active image area of a video signal. The line length is adjusted within the parameters specified in SMPTE 292M and 274M while keeping the 74 MHz HD clock and sync words in place in order to maintain a valid video signal. The vertical blanking within the video signal is also maintained for the storage of ancillary data. As such, HSDL can simply be thought of as two additional video formats. Figures 2 and 3 illustrate how the 2K data is placed within each video format.

2. daVinci, "How to Easily Move and Utilize Film Images as Data in a Real-World Video Facility."

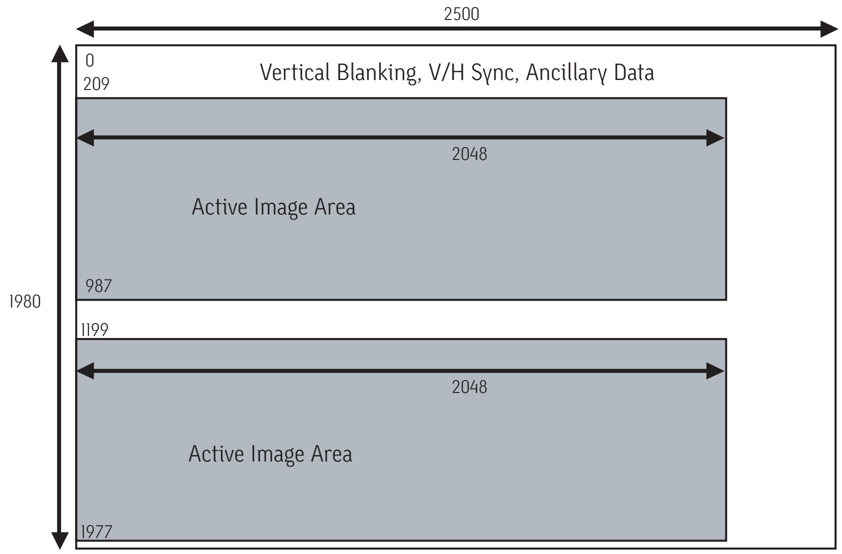

Fig. 2. HSDL/2K Data Frame at 15fps

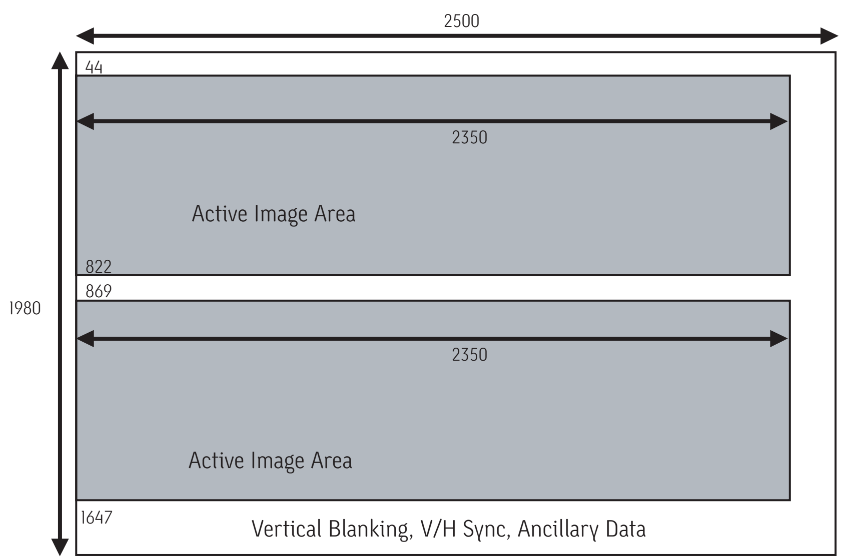

Fig. 3. HSDL/2K Data Frame at 19fps

The transfer of 2K RGB images with HSDL effectively uses only three-quarters of the 2.970Gb-per-second dual-link bandwidth. RGB is carried as 4:2:2 on link A, and 0:2:2 on link B with filler data passed for the 4:0:0 bytes on link B in order to maintain a legal video signal. As such, RGBA data could be transferred using HSDL without impacting the frame rate.

3.4 Fibre Channel

Fibre Channel technology is typically utilized within the production workflow to construct storage area networks (SANs) and directly connect disk arrays to computer systems. Supporting optimal transmission rates at up to 2.12Gb per second at distances up to 30 km, Fibre Channel can be used as a general-purpose network carrying ATM, TCP/IP, and other protocols but is more traditionally used to carry SCSI traffic between computer systems and disk arrays. Table 1 outlines the fibre channel topologies available.

Table 1. Fibre Channel Topology Configurations

Arbitrate Loop	Connects up to 127 nodes without a switch. All devices share the bandwidth, and only two can communicate with each other at the same time, with each node repeating the data to its adjacent node.
Switch Fabric	Connects computer systems and storage devices in a switch topology that permits all devices to directly communicate with each other. Multiple connections and multiple switches can be used to provide failover capabilities.
Point-to-Point	Direct connection between two fibre channel devices. Typically utilized to directly connect a storage device to a computer system.

With a single connection realized bandwidth of approximately 190MB per second, multiple Fibre Channel loops are utilized between devices to achieve the required total bandwidth required. For example, to achieve the more than 300MB-per-second bandwidth required for the real-time transfer of 2K film data, 2 2Gb-per-second Fibre Channel connections are required between devices. Next-generation Fibre Channel technology promises to deliver 4.25Gb per second. Facilities and equipment manufacturers searching for a cheaper alternative to Fibre Channel and the other technologies listed above look to Ethernet technology to implement data communications within the digital post-production workflow.

3.5 Ethernet

Since its invention in 1973, Ethernet (IEEE 802.3) has grown to meet the increasing bandwidth requirements of packet-switched networks. From 1Mb per second to 10Mb per second, 100Mb per second, and now 1,000Mb per second, also known as 1-Gigabit Ethernet (IEEE 802.3z), Ethernet has a lower implementation cost compared with other technologies and is reliable and relatively simple to install and maintain. Currently, copper and optical 1-Gigabit Network technology is being deployed in facilities, and the draft standard specifications for 10-Gigabit Ethernet (IEEE 802.3ae) are nearing completion. 10-Gigabit Ethernet will differ from earlier Ethernet standards in that it will support only optical fiber connections and operate only in full duplex mode, meaning that collision-detection protocols will no longer be necessary. 10-Gigabit Ethernet also extends the distance over which Ethernet can be utilized from the 5 km that Gigabit Ethernet supports to 40 km.

The advent of higher-bandwidth Ethernet technology has led to the creation of the Internet Small Computer Systems Interface (iSCSI), a protocol for transporting I/O block storage over an IP network. The Internet Engineering Task Force (IETF) is developing this protocol. In the past, the use of Ethernet limited the transfer of data to file granularity using FTP, CIFS/Samba, and NFS. With iSCSI, Ethernet can transport both standard network messaging and I/O storage blocks over the same wire. This ability permits SCSI data storage commands to travel over standard Ethernet connections. However, because iSCSI traffic could be traveling over a non-secure external IP network, security gateways or firewalls, as well as encryption, should be used to ensure security in a similar way that they would be used for normal network traffic.

iSCSI attempts to unify network-attached storage (NAS) and storage area networks (SANs) by using networking technology as a transport. Devices that support iSCSI contain an iSCSI host bus adapter (HBA). This HBA takes SCSI-3 data and encloses it in an IP packet that is then sent out via the Ethernet network. The receiving device strips off the IP packet information and passes the data to the SCSI controller. iSCSI is completely transparent: systems see what appears to be the SCSI controller while the network sees only IP traffic. iSCSI uses standard Ethernet infrastructure and standard SCSI protocols. When 10-Gigabit Ethernet becomes available, iSCSI will provide a SAN infrastructure that is five times faster than those built on currently available 2Gb-per-second Fibre Channel technology. iSCSI also permits storage traffic to move over a longer distance than the current 30 km limit of Fibre Channel. On the downside, because iSCSI requires the passing of IP and TCP headers, and checksums along with the SCSI commands, iSCSI data transfer has more overhead than equivalent communication over a Fibre Channel SAN.

Existing fibre channel devices can be accessed via iSCSI-to-Fibre Channel storage routers and switches. Additional protocols being developed by the IETF permit the use of TCP/IP with existing fibre channel SANs.

- FCIP encapsulates fibre channel packets and transports them via TCP/IP. A tunneling protocol, FCIP uses TCP/IP as the transport while keeping the Fibre Channel transport services in tact. This technique permits the interconnectivity of SANs across much longer distances.

- iFCP, on the other hand, connects fibre channel devices via IP networks by replacing the lower-layer Fibre Channel transport with TCP/IP and Ethernet transport. This permits IP-based SANs to link to both fibre channel devices and fibre channel SANs. With iFCP, Fibre Channel devices connect to an iFCP gateway or switch. At this switch each Fibre Channel session is terminated and converted to a TCP/IP session via iFCP. A second gateway or switch receives the iFCP session and initiates a fibre channel session with the receiving fibre channel device.

As the industry migrates from 2K data to 4K and even larger data resolutions, another alternative being examined by manufacturers to replace HIPPI and GSN, in addition to 10-Gigabit Ethernet, is Infiniband™.

3.6 Infiniband

Infiniband is a proposed new high-bandwidth switched fabric data communication standard that promises 2.5, 10, and 30Gb per second over copper wire at a maximum distance of 17 m and fibre optic cabling at 100 m to 10 km. Unlike earlier technologies that connect to a system using an existing bus interconnect such as PCI, Infiniband specifies a new host channel adapter (HCA) with a direct connection to the memory controller in the host computer system.

Infiniband can be used to directly attach two computer systems in a tightly coupled cluster configuration or as a switched-fabric network where HCAs in computer systems will connect through switches to target channel adapters (TCAs) in storage and other devices that support Infiniband. As in the case of other switched-fabric technologies, reliability is achieved through redundant fabric links and automatic failover in the switches. Switches that connect Infiniband to existing fibre channel and Ethernet networks are being developed in order to ease the transition to this new technology. Trials of this new interconnect technology are now complete, with deployment scheduled to begin in 2003.

The above technologies are positioned to support high-bandwidth data communication within a production facility. When large amounts of data must travel between remote facilities over a long distance, SONET is typically utilized.

3.7 SONET

The Synchronous Optical Network, or SONET (B-ISDN), is a fibre-optic transmission system utilized by telecommunication companies to provide high-speed digital services between production facilities. SONET uses time division multiplexing (TDM) to send multiple data streams simultaneously at speeds ranging from 51Mb per second to 39Gb per second. Table 2 outlines the exact levels of service available.

Table 2. SONET Levels of Service

Service	Speed (Mb/sec)
OC-1	51.84
OC-3	155.52
OC-12	622.08
OC-48	2,488.32
OC-192	9,952.28
OC-768	39,813.12

To ensure reliability, SONET is built in a self-healing ring architecture that requires at least two transmission paths in the event that one fails. Asynchronous Transfer Mode (ATM) technology is typically used to create a logical end-to-end circuit that guarantees quality of service.

4.0 Digital Archive

The hub of the cinema mastering workflow is the central digital archive. This is the repository of all the digital assets for the current production. Here, all data files, along with all associated metadata for a production, are stored in a database for use during post-production. As required by post-production, assets are checked in, checked out or locked. They are then modified, and resubmitted as a new version with all changes tracked in the metadata so that they can be backed out in order to restore a previous version if necessary. The data archive must also provide efficient content management. As such, assets within the digital archive must be searchable such that they may be called up and viewed at any time. The archive must also support the creation of proxy images and movies for browsing and the instant preview of sequences on desktop and laptop preview stations.

By storing data in a central archive the dataflow is no longer constricted linearly even if the overall workflow remains linear. Instead of the data storage moving from one system to another as it is required, with the possibility of data loss at each step, the data is stored in one central location. This permits data access by all systems, allowing the potential overlap of work, and the potential to go back and repeat an earlier step in the process without physically moving the data upstream. The central data archive also enables the efficient backup of all data from one central location.

Three basic strategies exist for the implementation of a digital archive. These include a client/server approach using direct-attach storage (DAS), network-attach storage (NAS), and storage area network (SAN) technology. A SAN with support for a common filesystem among all systems provides a scalable resilient digital archive solution.

4.1 Storage Area Network

In order to effectively share the digital assets of a production within a facility, the disk storage within the central data archive should compose a storage area network. In a SAN (figure 6), storage devices, as well as computer systems, are directly connected to a high-speed switched-network fabric that permits all attached devices to communicate directly with one another using a block data access protocol such as SCSI. This architecture is in contrast to DAS, such as that in figure 4, where storage devices are directly connected to computer workstations and servers. This topology creates islands of information on disparate systems. SANs eliminate these islands of information, because each storage device is available to each computer system on the network. This eliminates the requirement that data be copied between workstations.

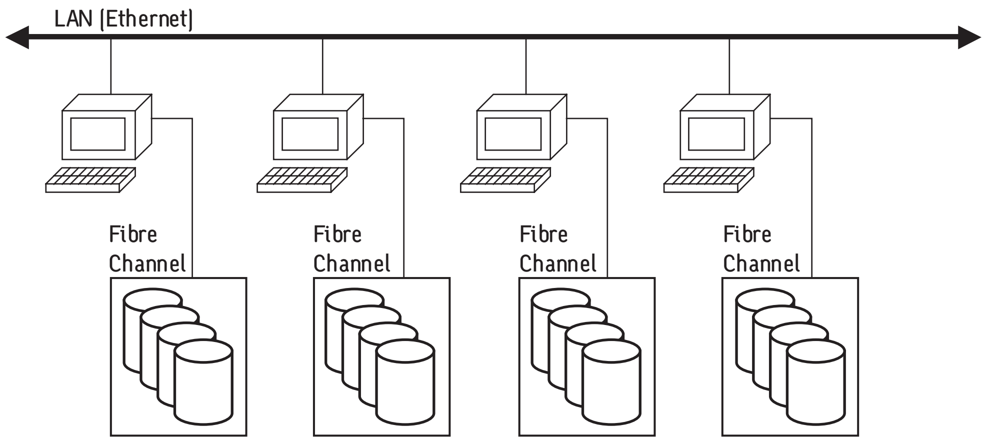

Fig. 4. Direct-Attach Storage (DAS)

Today this high-speed network fabric is typically composed of interconnected Fibre Channel switches. These switches provide point-to-point connectivity between computer systems and storage devices. Multiple Fibre Channel links are utilized to achieve the bandwidth requirements of 2K image data transfers. In the case of 1Gb per second Fibre Channel, four links are needed between computer systems, storage components, and switches. With 2Gb per second Fibre Channel technology, the number of required links is reduced to two. To date, SANs that deliver 380MB per second with 1Gb-per-second Fibre Channel and 800MB per second with 2Gb Fibre Channel have been deployed to support the cinema-mastering workflow. Longer term, the availability of 10-Gigabit Ethernet technology will make SANs based on TCP/IP network technology possible.

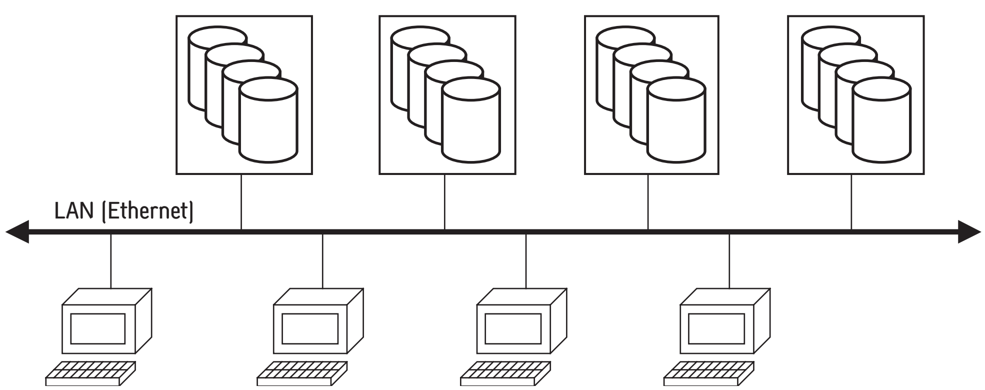

Fig. 5. Network-Attach Storage (NAS)

In addition to being an improvement over a DAS architecture, a SAN is also an improvement over more traditional network-attach storage architectures (see figure 5), which utilizes file access protocols like CIFS/Samba and NFS transported over Ethernet and TCP/IP to provide file access to client systems. The SCSI block data access protocol used in a SAN eliminates the overhead required for CIFS and NFS and as a result provides higher bandwidth to applications. Block data access also provides for higher granularity access than the file granularity provided by CIFS and NFS. For example, whereas Ethernet transfers are limited to 1512 bytes per packet, Fibre Channel technology, on which current SAN infrastructure is based, permits an application to negotiate the transfer window size. This size can vary from a few bytes up to 200KB. This ability to handle large bulk data transfers with low overhead lends itself well to the requirements of the cinema-mastering workflow, which requires the transfer of sequences of large image files and ancillary data between storage devices and computer systems that implement the post-production workflow.

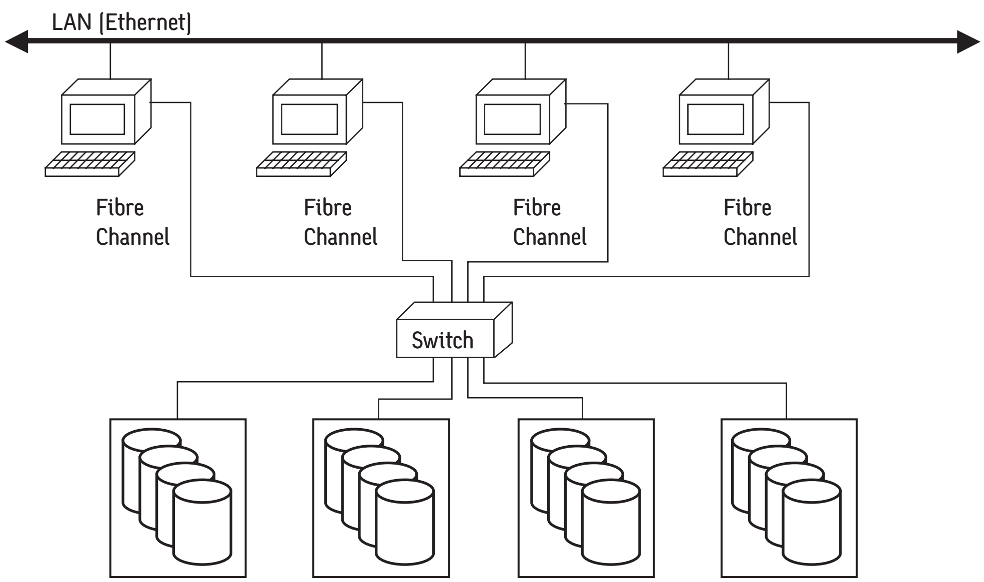

Fig. 6. Storage Area Network (SAN)

Utilizing a SAN as the basis for the digital archive infrastructure also provides other benefits to the digital cinema mastering workflow. These additional benefits are outlined in table 3.

Table 3. Storage Area Network Benefits

Scalability	Storage can be added as it is required by a production. So, although a typical production might require 3TB of storage, this storage does not all have to be available at the start of production. Storage might be added as it is no longer needed by another production that is ramping down, thus allowing multiple projects to utilize the shared storage most efficiently.
LAN-Independent Backup	Because data passes over the higher-performance Fibre Channel network, data in the digital archive can be backed up without flooding the LAN. As such, the performance of LAN-based applications is not impacted.
Lower Cost of Ownership	Because storage devices attach directly to the network, server systems are not required for storage attachment. This reduces both hardware cost and administrative costs.
Resiliency	Redundant connections to switches, combined with automatic failover support in switches, provide a robust network with high availability.
Data Recovery	RAID storage devices with hot swappable drive components permit the replacement of a failed drive and the automatic recovery of data without the need to restore the content from backup media.

One other distinction that must be made is that between a more traditional SAN and a digital production SAN. In a traditional SAN, client workstations are attached via a client-access LAN. In this traditional model, applications run on larger servers that participate in the SAN and subsequently serve lower-bandwidth clients over the client-access LAN. In order to meet the higher-bandwidth needs of post-production, the workstations running the software applications required for post-production participate directly in the SAN. Since these workstations typically have different system architectures, a common filesystem is required. And, to ensure high data availability, the workstations and servers within the workflow are managed as a cluster.

4.2 Common Filesystem

Simply implementing the data archive as a SAN is not enough to meet the requirements of the data-centric digital production workflow. Today, the software applications that compose the tools within this workflow are implemented within a production facility by computer systems with many different system characteristics. In addition to differences in the operating system, these differences include 32- versus 64-bit addressing and big endian versus little endian issues. As such, the underlying filesystem on which the central data archive is built must support the connectivity of these heterogeneous clients inorder to promote the efficient sharing of data within a facility and demonstrate a uniform view of the data at all workstations within a facility. A common filesystem provides a consistent view of the assets within the archive and permits access to these assets from any system on the SAN.

This common filesystem should be implemented at the operating system kernel level and provide for binary file compatibility on all systems within the SAN. This is in contrast to efforts such as the Advanced Authoring Format (AAF),[3] the goal of which is to provide file compatibility between software applications used within the workflow. Whereas AAF provides a common file format for the sharing of data and associated metadata between different software applications, the benefit of a common filesystem in a SAN is to provide for the transparent sharing of files between different computer systems that compose the SAN. As a result, the common shared filesystem ultimately makes a SAN much more useful within the digital production workflow. For example, once the high-resolution images for a scene have been ingested into the SAN, the data is instantly available to the technician doing dust busting and scratch removal, without the requirement that the data be copied to the system on which the dust-busting process will take place. When dust busting is complete, the data is instantly available for the next process, such as color grading.

Several common filesystem architectures have been proposed and implemented. One implementation is the CXFS™ filesystem from SGI. CXFS[4] is a 64-bit journaled filesystem that utilizes asynchronous data-buffering techniques to avoid unnecessary physical I/O operations. This practice permits the efficient allocation of disk space into the large contiguous chunks required by the post-production workflow. Also, because the disk space is allocated in large contiguous blocks, high-bandwidth transfers can be sustained. A token-based scheme is used to control concurrent file access. In the event of a system or disk failure, inspecting the journal and rolling back any incomplete transactions achieve data recovery. File data flows directly between systems and disk over the SAN, while filesystem metadata travels between systems using a secondary 100Base-T or Gigabit Ethernet network. This metadata contains information that describes each of the files in the filesystem, as well as information about the filesystem itself. This metadata keeps each of the systems up to date on the state of the common filesystem. In order to ensure the integrity of this metadata and to ensure high availability of the data for a production, the systems within the workflow that participate in the SAN are managed as a cluster.

4.3 Cluster Management

The management of the systems that compose the SAN and digital archive as a computer cluster eases system administration tasks while at the same time providing redundancy and failover capability to ensure resilient access to the data stored within the archive. Cluster management permits systems to be added to and removed from the SAN from a single control console. The mounting and unmounting of storage volumes, as well as the setting of file access control parameters, can also be performed from this single administration console.

5.0 Security

In the entertainment industry intellectual property takes the form of the digital assets stored within the data archive. As such, security is a prime concern throughout the workflow. This security takes two forms. The first is access control to ensure that only those personnel who require access to the data in order to complete their work on a production have the level of access that they require. The second is data protection in the event of a component failure.

3. *http://www.aafassociation.org/*

4. SGI White Paper, "SGI CXFS: A High-Performance, Multi-OS SAN Filesystem from SGI."

5.1 Access Control

This security is even of greater concern when the digital workflow connects multiple facilities. This security takes the form of encryption and access control technologies. Various techniques exist to guarantee the security of the data throughout the workflow.

- Individual user accounts with expiring passwords should be maintained for all personnel who require access to production data.
- Computer systems within post-production should automatically log off users after a short amount of idle time. This practice prevents an idle workstation from being used and data being accessed by an unauthorized user.
- Internal IP networks should be isolated from public networks by firewalls that filter out unwanted data access from the outside.
- Encryption should be used when data is transported over a public network. Encryption technologies can also be used to encode backup media to prevent unauthorized access.

5.2 Data Protection

As previously discussed, the first step in the protection of the digital assets for a production is the use of a digital archive. This archive centralizes data storage in a single place and eliminates multiple copies and islands of information on disparate systems. Once the data is stored in a digital archive, a blend of techniques provides comprehensive data protection.

- Redundant array of independent disks (RAID), with the ability to recover from the failure of a physical disk device within a disk enclosure, prevent data loss at the lowest level.
- Multiple connections between systems, switches, and storage with failover capabilities provide alternative paths to data in the case of a primary path failure.
- Snapshot, or point-in-time, data copies create an image of data on disk at a particular moment in time to permit fast recovery in the event of data corruption.
- Mirroring, or the creation of concurrent copies of critical data both to a local storage archive and to a remote location, protect against hardware or site failures.
- Backups to either tape, or disk, or both protect against data corruption, hardware failures, and user errors. Tape backups, while slower to create and restore, enable offsite storage. Disk backups are faster and offer the ability to quickly create a disk image that can be backed up to slower nondisk media without impacting ongoing operations.
- Hierarchical storage management (HSM) systems optimize storage device usage by transparently moving less-frequently-used data to secondary storage, which is backed up less frequently. This reduces the time required to backup the primary storage medium.

6.0 Archival Storage

After the production of a feature film is complete, all the digital assets for a film must be stored in such a way that they can be located and easily reloaded into the workflow. This process would be required for the production of a sequel, to utilize data from one production in a later production, or to re-master the original content for additional copies or into additional distribution formats. In the past, archival storage took the form of a final film print. However, taking a completely digital master and archiving it on celluloid is not an optimal solution in the digital workflow. Once archived on film, the production of a sequel requires the scanning or redigitizing of the archival film print. This procedure adds another generation of data loss to what was once already digital material. Another factor to consider is that, along with the final master, there is metadata and other data files, such as those used to create special effects models and the like. This type of information cannot be archived on film. Storing this information along with the final images on a digital medium keeps all information for a production together for reuse at a later date. And while it is true that, as with celluloid, most digital media also have a limited shelf life before deterioration begins, material stored in a digital medium can be moved to a new or next-generation digital medium without data loss, whereas copying a film print or retrieving data from an archival film print will always result in a generation of data loss.

Archival storage is either nearline or offline. Shortly after a production is complete and a film is released, the data is kept in nearline storage to be used for additional copies, or for a sequel or other promotional materials. This nearline storage typically takes the form of digital tapes within a tape robot system directly connected to the SAN. Longer term, these tapes are removed from the tape robot and stored offline in a temperature-controlled secure vault. This makes room in the tape robot system for additional tapes to hold more recent productions. Digital assets moved to the vault should be recorded in a searchable database that facilitates the retrieval and reloading of digital assets into the workflow as they are required.

7.0 The Complete Picture

The combination of a digital archive built around a SAN with a common filesystem with GSN and HSDL utilized between non-SAN-capable devices provides an ideal digital cinema mastering workflow implementation. In this sample implementation (see figure 5), the data enters the archive by way of a Thomson Spirit DataCine™ or a Thomson Viper FilmStream Camera™.

From the Spirit DataCine system, the 2K data passes over a GSN interconnect at 24fps to a computer system, where it is buffered on local disks. From this system the data is then transferred to the digital archive via a secondary data interface that connects it to the SAN. This secondary interface requires 2 parallel 2Gb per second Fibre Channel links to provide adequate bandwidth to transfer the frames in real time.

From the Viper Camera system, the 10-bit RGB data passes over a dual-link SMPTE 292M connection to the local storage on a host computer. Because shooting is typically done at a remote location, the data is later transferred to the digital archive by attaching the disk storage to a SAN-attached computer system. In this case, the data is then transferred over fibre channel data links to the central digital archive.

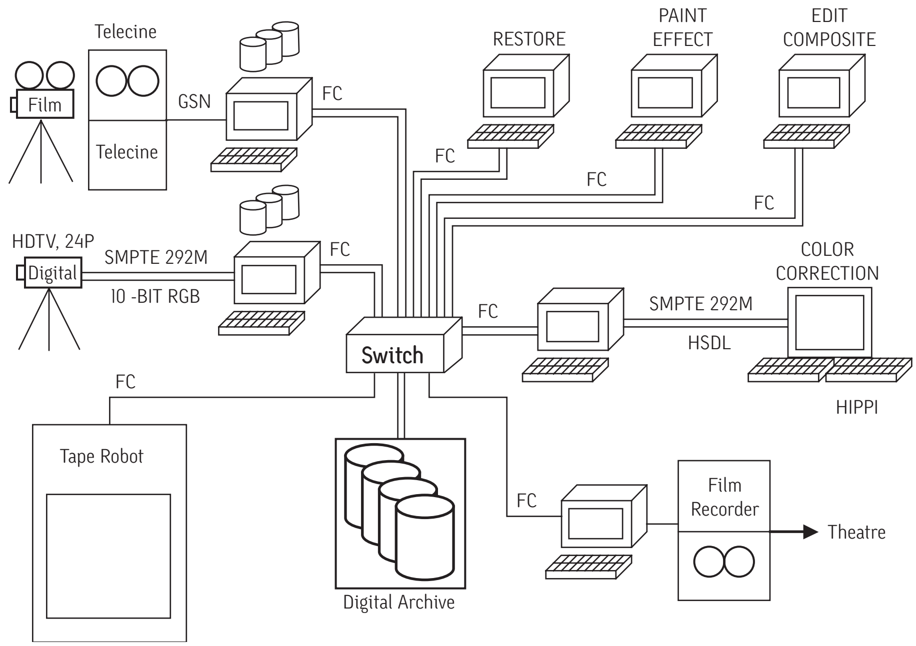

Fig. 7. A Data-Centric Cinema Mastering System Implementation

Post-production work includes restoration (dust busting and scratch removal), as well as editing, compositing, and paint and special effects work performed on one or more computer workstations directly connected to the SAN via one or more high-speed fibre channel data links. For color grading, image sequences are transferred to and from the da Vinci 2K™ color corrector using HSDL over a dual-link SMPTE 292M connection. Since the color corrector is a special-purpose system that is not SAN aware, a host computer with a high-speed connection to the data archive sits on the SAN and feeds image data to and receives data back via HSDL from the color corrector. Color-graded frames are then saved back in the archive.

For film distribution, the final print is realized with an ARRILASERTM film recorder connected to the SAN via a host computer. This system grabs frames from the digital archive and passes them to the laser-recording device in non-real time. Finally, the data is archived directly onto DLT tapes using a StorageTek® tape robot connected directly to the SAN fabric. This tape library can also be used to create the backup tapes for offsite storage.

8.0 Remaining Challenges

This paper describes network, data communication, and storage technology that enables a completely digital cinema mastering workflow from acquisition, through post-production, to distribution from a single universal digital master. However, significant challenges remain that have not been described here.

- Asset Management: With several productions ongoing within a facility at any one time, there is a significant need for a robust digital asset management system that tracks all assets within the archive, as well as those that have been archived to offline storage, and permits searching and browsing of all assets. Database software that supports these features exists, but this software is general purpose and not specifically tailored to the requirements of the cinema mastering workflow.
- Metadata: Another challenge is the comprehensive system handling of metadata for all assets. All changes

made to assets must be recorded in the metadata that accompanies each data file. Efforts are ongoing within the AAF Association to address this issue.

- Guaranteed Rate I/O: Switched fabric networks fail to provide the guaranteed I/O bandwidth found between direct attach devices. Although sufficient bandwidth is available between a device and a switch, the switch becomes a bottleneck as network traffic increases. Switches provide the capability to lock out connected devices and concentrate bandwidth on particular connections, but advances in switch technology are required to enable smart load balancing.

9.0 Conclusion

The cinema mastering workflow, once the domain of processing labs and the negative cutter, has been revolutionized by digital technology. The ability to capture, store, and transmit high-resolution digital images in real time has revolutionized feature film production while at the same time making it possible to create several distribution formats from a single universal master. This revolution has refocused the production workflow on the data within the workflow and, as a result, has significantly changed the technical requirements of the process. Today, production facilities must implement high-bandwidth communications, flexible storage, secure data transfer, and digital archive capability in order to effectively realize the benefits of this new methodology.

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America (52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 6771.0290

3428 (12/16/2002) J14185