

High-speed storage for fast-paced studios

From large production studios to two-person boutiques, the mantra is the same: More bandwidth and more capacity.

BY DAVE SIMPSON

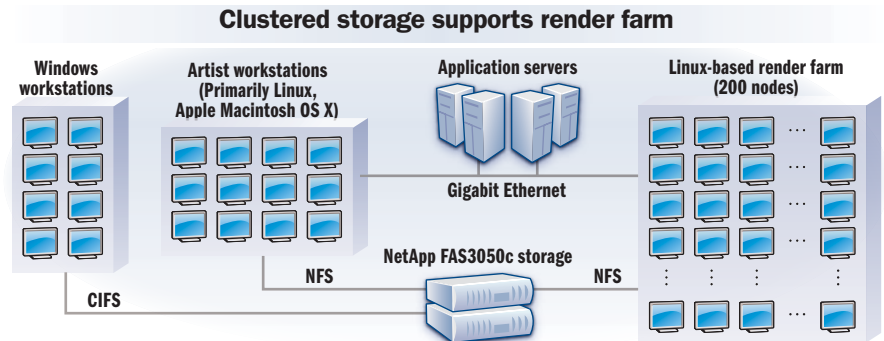
The storage requirements for media/entertainment facilities are unique: blazing bandwidth, capacious capacity, and a need for collaborative content sharing, as illustrated in the following case studies:

SPIN WEST VFX

SPIN West VFX, a franchise of SPIN Productions in Vancouver, BC, is used to juggling a variety of TV, film, and commercial projects simultaneously. The studio specializes in visual effects, animation and design work. Recent film projects included *Land of the Dead* and *Outlander*, and the studio has more than 40 artists and composers.

The studio's storage capacity and bandwidth requirements are massive. For example, one particular shot in *Outlander* required almost 1TB of assets. And there were sometimes 1,000 frames per shot.

SPIN West VFX has more than 200 Linux-based render nodes, often running 24x7, hitting its storage systems as artists try to pull off the same data to render out particular frames. Previously, the studio used dual-controller "white box" storage hardware with clustered file system software. But that storage setup created performance bottlenecks that sometimes resulted in time outs during heavy rendering, in which case the artists would have to re-start the rendering



SPIN West VFX uses a NetApp FAS3050 clustered storage system to support a 200-node render farm and more than 40 artists.

from scratch.

SPIN West VFX's IT department solved the problem with a NetApp FAS3050C running the Data ONTAP GX operating system. The configuration provides a clustered storage system with global namespace technology that allows IT to maintain an existing logical file directory structure, despite additions and changes to the underlying hardware. SPIN West also takes advantage of NetApp's FlexVol functionality, which is used for load balancing to maximize performance with shifting rendering requirements.

With the global namespace technology, the studio can put files anywhere within the cluster, and can add more controllers when it needs more horsepower. IT administrators at SPIN West expect the initial 10TB capacity to grow to about 20TB over the next year.

SPIN West VFX currently uses SATA drives in the FAS3050C, in a RAID-DP (RAID 6) configuration for added reliability, but the studio plans to add Fibre Channel drives for increased performance and use the SATA disks as near-line storage.

The key advantages of implementing the NetApp FAS3050C with Data ONTAP GX, compared to the previous storage configuration,


included:

- Improved reliability (no downtime);
- Easy scaling to add more storage controllers without disrupting production;
- Rapid load balancing with FlexVol software;
- Ability to re-use "legacy" hardware.

SANTA FE PRODUCTIONS

When Santa Fe Productions, a post-production company in Brooklyn, NY, hit the storage capacity/bandwidth wall, the studio investigated a Fibre Channel SAN and 10GbE-based storage solutions. But those options proved to be twice as expensive as the storage configuration that Santa Fe wound up implementing, which includes Atto Technology's FastStream SC 7500 controller and 4Gbps Celerity HBAs, as well as UltraStor RS16JS SAS/SATA disk arrays from Enhance Technology.

Before upgrading its storage environment, Santa Fe Productions had two Mac workstations and was swapping HD files between direct-attached eSATA/FireWire disk arrays. That setup proved time-consuming and interrupted the production studio's workflow, according to director and editor Eli Brown. And when the firm added a third Mac, data



The storage configuration behind the creation of images for *The Brain Fitness Program* included 4Gbps Fibre Channel HBAs and a FastStream SC 7500 RAID controller from Atto, and SAS/SATA JBOD disk arrays from Enhance Technology.

in an all-SAS configuration). And the controller provides the ability to connect up to four workstations and to add up to eight 16-bay arrays for more capacity.

Essentially, Atto's SC 7500 provides the shared-storage environment and features of a SAN, but without the expense associated with a traditional switch-based Fibre Channel SAN, according to Brown. Other features of the SC 7500 include built-in RAID functionality and four external Fibre Channel ports

Santa Fe Productions creates non-fiction programming for broadcast television and home video and film distribution. One of its latest projects was *The Brain Fitness Program*, which aired on PBS stations.

FRAMESTORE

One thing's for sure when it comes to storage in the media/entertainment industry: Capacity grows rapidly. For example, in creating graphics for the feature film *The Tale of Despereaux* (from Universal Studios), UK-based Framestore started with 50TB of capacity and by the end of production the studio had a whopping 200TB (more than 90 million files)—400TB if you include a mirrored cluster. All of the animation and computer-generated imagery (CGI) content was stored on Infortrend's EonStor A24F-R2430 SAS/SATA RAID arrays in a RAID-6 configuration for added reliability. (RAID 6, or dual-parity RAID, protects data in the event of simultaneous drive failures.) During peak production, Framestore was generating 5TB of animation content per day.

Framestore uses SAS drives for metadata storage and SATA drives for bulk storage. The EonStor RAID arrays are configured with external Fibre Channel interfaces.

The CGI studio created a 200TB storage cluster using the Linux-based Lustre open-source cluster file system. The cluster was tightly coupled to a 6,000-core render farm.

According to Steve Prescott, Framestore's technology director, the studio is getting performance of about 3,000MBps and 50,000 I/Os per second (IOPS) on the cluster-storage configuration.

For data protection, Framestore used mirroring and snapshots (in addition to RAID 6) over a 10Gbps dark fiber network installed between its three Soho production offices and its data center in London.

The *Tale of Despereaux* was the first project for Framestore's feature animation division. Fun facts:

- There are 40,089 individual assets in the film;
- 1,726 final shots were delivered;
- The film has 126,248 frames;
- The animation crew size was 278 during peak production;
- And...there are 413,138 hairs on Despereaux's head.

Some of Framestore's other recent projects include *The Chronicles of Narnia: Prince Caspian*, *The Dark Knight*, *Quantum of Solace*, and *Australia*.

ADVANCED DIGITAL SERVICES

Hollywood-based Advanced Digital Services (ADS) started out as a tape duplication facility, but has since branched out into a post-production and media services company that does work such as digital restoration, encoding, audio restoration, digital transport of content and other services—mostly for TV.

On the storage front, one of ADS' configurations consisted of five non-linear editing workstations (mostly Final Cut Pros) attached to an Apple Xsan SAN configuration, but some of the studio's editing stations needed to deal with bandwidth requirements that the Xsan couldn't handle (including 440Mbps data rates from Sony SRW systems), according to Russell Ruggieri, ADS' director of engineering.

To eliminate those bandwidth constraints on the workstations that needed the highest throughput rates, ADS installed two BlueStor PeSAN RAID arrays from JMR (for a total capacity of 128TB on SATA disks), which are unique because they use the PCIe bus to directly connect systems to BlueStor disk

rates dropped precipitously.

"We had to duplicate all of our digitized footage and other graphic assets across multiple FireWire drives and share them as we needed them," Brown explains. "If someone needed a clip and they didn't have it, we'd have to copy it from wherever it was and move the drives around."

Santa Fe Productions' current storage configuration is anchored by Atto's FastStream SC 7500 controller, which sits between the three Mac workstations (a G5 and two Mac Pros running a variety of video applications, as well as Tiger Technology's MetaSAN 3.0 SAN management software) and 16-bay UltraStor RS16JS SAS/SATA JBOD arrays from Enhance. The total usable capacity is 13.7TB, with 1TB SATA drives in a RAID-5 configuration.

The new storage environment allows the studio to run multiple compressed HD video streams simultaneously without any bottlenecks. "The key benefit of the SC 7500 was in workflow productivity," says Brown. "We have more than enough performance, and we can share files without losing time copying them." (Atto claims that the SC 7500 can support up to five streams of 10-bit uncompressed HD video, with performance of up to 1,200MBps



Russell Ruggieri, Advanced Digital Services' director of engineering, uses BlueStor PeSAN RAID arrays, which leverage the PCIe bus for host-storage connections.

arrays. (In the case of ADS, PCIe is used for connecting to Mac platforms while PCI-X is used for connecting to Windows platforms.)

Ruggieri says the new storage configuration solved his performance problems, providing more than 1,000MBps of throughput and the ability to handle 4K-resolution formats. "It also makes it easy to add capacity, and JMR's PCI adapters give me additional slots on the Macs," says Ruggieri. "And it was relatively low cost." (JMR cites pricing of less than \$1 per GB.) ADS may replace the remainder of its 2Gbps Fibre Channel configurations with additional BlueStor systems.

JMR's BlueStor PeSAN (PCIe SAN) RAID arrays come in a 3U, 16-drive enclosure with up to 16TB with SATA drives or 4.8TB with higher-performance, 15,000rpm SAS drives. Because they use a switched architecture (with PCIe X4 or X8 internal switches) and the PCIe bus for host-storage interconnection, the arrays eliminate overhead associated with Fibre Channel, Infiniband or Ethernet protocol conversion. JMR claims performance of more than 2,000MBps. Targeted applications include content creation, SD/HD video editing, and 2K/4K digital intermediate applications.

HOLLYWOOD INTERMEDIATE

Hollywood Intermediate has provided post-production and digital intermediate (film to digital to film) services on more than 70 motion pictures. The company is also performing post-production work for a network television series.

The current storage capacity at Hollywood Intermediate (HI) is about 100TB, and the studio is in the process of nearly doubling that capacity. The facility has a mix of 2Gbps Fibre Channel systems and some 4Gbps Fibre Channel disk arrays; all of the new storage devices will be 4Gbps Fibre Channel, and the 2Gbps Fibre Channel arrays will be relegated to proxy storage and final output storage, according to David Graubard, a senior systems engineer and architect at HI.

For its primary storage, Hollywood Intermediate uses 4Gbps Fibre Channel disk arrays from Arena Maxtronic, including dual-controller 16-bay SS-6601R and 24-bay SS8801R arrays, single-controller SA6651Es, and some older 2Gbps Fibre Channel models. According to Graubard, the key purchasing criteria were price, configuration flexibility, capacity, and performance.

"A frame of film scanned at 2K resolution is about 12MB, which is manageable, but we may need 24 of those 12MB frames to make up one second, which could require 292MBps of throughput—that's more than twice the throughput needed for uncompressed HD video," says Graubard. "For the average 2-hour movie, we may need a minimum of 2TB of storage, but since we scan in the entire segment of a scene, even if they use two seconds of that scene, an average movie can chew up about 12TB to 16TB on our hard drives. Since we are typically working on quite a few movies at once, this can easily reach into the hundreds of terabytes of storage. Of course, many directors want their movies scanned in at 4K resolution these days, so multiply all those numbers by four. That's about 1.2GBps of throughput per machine working in that resolution."

All of Hollywood Intermediate's Arena Maxtronic arrays are configured with SATA drives (although some of the arrays also support SAS). "SAS is very fast, particularly for small files and databases, but our smallest file format is about 8MB per frame, and we run at 24 frames per second, and SAS doesn't provide a lot of performance improvement in that scenario," says Graubard.

AN IDEAL WORLD

For creative professionals on the move, small, lightweight, portable RAID arrays with ▶ p. 32



To create HD visual effects for many of its commercials, such as this one for Korbel Champagne, An Ideal World uses a Fusion F2 RAID array from Sonnet Technologies.

high performance are the way to go. Those are some of the reasons that Robb Hart, co-founder of An Ideal World (a production/post-production studio in Santa Ana, CA) opted for Sonnet Technologies' Fusion F2 disk array.

An all-Mac shop, An Ideal World specializes in visual effects, primarily for TV commercials. Although a small studio, An Ideal World uses some cutting-edge equipment, including high-speed Phantom HD cameras and 4K-resolution RED ONE cameras. Some of the firm's recent projects included graphics for Korbel Champagne Cellars and Barona Casino commercials.

"The 4K RED camera shoots very high-res images digitally [as opposed to film], but that leads to a lot of data," says Hart, who uses Sonnet's Fusion F2 connected to MacBook Pro laptops via eSATA connections. "I needed a storage device with a lot of capacity and performance, and one that was portable because we move around a lot from location to location." Hart also likes the fact that the Fusion F2 does not require an external power supply. (The array draws its power from the laptop's FireWire port or an optional power adapter.) "With the high-res cameras I'm using, there was no way I wanted to play back the shoot in SD," Hart says,

describing a recent commercial project. "I knew that on-set, HD playback would be possible if I could find the right RAID device."

For a recent commercial, Hart used the Fusion F2 RAID array with an AJA IoHD device to capture HD footage, recorded in Apple's ProRes codec, from the RED camera as a 720p (1280x720 pixels) downconversion and 1920x1080 Phantom material—both played back at a 29.97 frames-per-second (fps) rate. "Using SATA drives rather than FireWire drives enables playback of shots in HD," says Hart.

The studio's Fusion F2 array includes two 2.5-inch, 5,400rpm SATA drives with a total capacity of 640GB. (Fusion arrays can also be configured with 1TB SATA drives.) The system measures 5.9x6.2x0.72 inches. The drives are shock-isolated to minimize or eliminate cross-coupled vibration that can lead to soft errors, which are potential sources of skipped frames in video capture and playback. Sonnet claims performance of up to 153MBps on write operations, and 143MBps on reads (almost twice the performance of FireWire drives). ©

VENDORS MENTIONED

Apple, Arena Maxtronic, Atto Technology, Enhance Technology, Infortrend, JMR, NetApp, Sonnet Technologies, Tiger Technology.

Do you really need a SAN? continued from page 31

other key storage features. While many of the storage features within VMware leverage SAN array capabilities, the control and management is substantially shifted to the application realm.

Painful change?

Clearly, the concept of application-managed storage represents a significant shift from current understanding of maximum efficiency and capability in the enterprise, and a move to such a model would not be without pain. Enterprise storage changes very slowly, with new architectures having to fit refresh cycles to avoid disruptive and costly forklift upgrades. IT purchasers are notoriously conservative, so significant momentum will need to be present before a majority will move in this direction.

Most significantly, the capabilities of the applications will need to continue to improve,

because such a shift cannot happen at the expense of application performance or data availability. In the current economic times, it will be incumbent on both pure storage vendors—and application vendors that may unseat them—that their solutions are the most effective choice, both operationally and economically.

Time will tell whether this becomes a true sea change in the enterprise, or if it is just a minor blip in the current evolution of enterprise storage. ©

ANDREW REICHMAN is a senior analyst at Forrester Research. To receive free related research from Forrester, visit: www.forrester.com/sanarticle

VENDORS MENTIONED

Hewlett-Packard, Microsoft, Oracle, VMware.

GROUP PUBLISHER

SHANNON ALO-MENDOZA
 98 Spit Brook Road, Nashua, NH 03062-5737
 Tel: (603) 891-9447 Fax: (603) 891-0587
shannona@pennwell.com

NATIONAL ACCOUNTS MANAGER

CAROL STAGG
 P.O. Box 9235, Calabasas, CA 91372
 Tel: (818) 348-1240 Mobile: (818) 424-9736
 Fax: (818) 348-1742
clstagg@aol.com

DIRECTOR, LIST SALES

BOB DROMGOOLE
 Tel: (603) 891-9128 Fax: (603) 891-9341
bobd@pennwell.com

UK, SCANDINAVIA

TONY HILL
 PennWell Publishing UK Ltd.
 Warlies Park House, Horseshoe Hill
 Uphire, Essex EN9 3SR UK
 Tel/Fax: +44 (0) 1442 239 547
tonyh@pennwell.com

FRANCE, NETHERLANDS, BELGIUM, S. SWITZERLAND, SPAIN, PORTUGAL

LUIS MATUTANO
 Adecome+, 1, rue Jean Carasso
 F-95870 BEZONS, France
 Tel: (33) 1 30 76 55 43 Fax: (33) 1 30 76 55 47
luism@pennwell.com

GERMANY, AUSTRIA, N. SWITZERLAND, E. EUROPE

HOLGER GERISCH
 Verlagsbuero Gerisch Hauptstrasse 16,
 D-82402 Seeshaupt, Germany
 Tel: (49) 8801-302430 Fax: (49) 8801-913220
holgerg@pennwell.com

JAPAN

MANAMI KONISHI
 e.x. press Co. Ltd., Aihara Bldg. 4F,
 2-13-1, Hirakawa-cho, Chiyoda-ku
 Tokyo 102-0093, Japan
 Tel: (81) 3 3556-1575 Fax: (81) 3 3556-1576
manami.konishi@ex-press.jp

HONG KONG/CHINA

ADONIS MAK
 ACT International
 Room 3401, 34/F, Hopewell Centre,
 183 Queen's Road East, Wanchai, Hong Kong
 Tel: (852) 2-838-6298 Fax: (852) 2-838-2766
adonism@actintl.com.hk

ISRAEL

DAN ARONOVIC
 Allstar Media Inc.
 3/1 Hatavas Street, Kadima 60920, Israel
 Tel/Fax: (972) 9-899-5813
aronovic@actcom.co.il

EDITORIAL REPRINT

ASHLEY GOLDSTEIN
 Tel: (603) 891-9114
 Fax: (603) 891-0597
ashleyg@pennwell.com

MARKETING COMMUNICATIONS

JONI MONTEMAGNO
 Tel: (603) 891-9110
 Fax: (603) 891-0624
jonim@pennwell.com

AD TRAFFIC MANAGER

DANIEL GREENE
 Tel: (918) 831-9401
 Fax: (918) 831-9153
danielg@pennwell.com

FOR SUBSCRIPTION INQUIRIES ONLY:

Tel: (847) 559-7500
 Fax: (847) 291-4816
ifs@omeda.com