



**INFOSTOR**

# STORAGE IN THE *Studio*

*A special section  
addressing the  
storage  
requirements of  
digital content  
professionals*

## INSIDE

### **Serial ATA means low-cost storage**

For many graphics applications, disk arrays based on SATA disk drives may provide as much performance as Fibre Channel disk arrays—at half the cost.

### **Studios move to network storage**

In media environments, there's a clear trend away from direct-attached storage to network storage, which provides a number of advantages for collaborative projects.

### **DAMming the flow**

Digital asset management (DAM) tools help digital content creation (DCC) professionals control rivers of data.

# Serial ATA means low-cost storage

*For many graphics applications, disk arrays based on SATA disk drives may provide as much performance as Fibre Channel disk arrays—at half the cost.*

**By Dave Simpson**

**D**igital content professionals have unique requirements when it comes to data and content storage. They need extremely high capacity and as much speed as they can afford. Unfortunately, most studios operate under tight budgets when it comes to capital expenditures for computing equipment.

A relatively new disk drive interface—called Serial ATA, or SATA—promises to solve those problems by providing significantly lower-cost storage subsystems than those based on disk interfaces such as Fibre Channel and SCSI, at performance levels that are “good enough” for most digital content applications, including animation, 3D rendering, non-linear video editing, and a variety of other rich media applications.

Choosing the right storage system for your environment comes down to a tradeoff between cost and performance. Although pricing on disk arrays varies widely, Charlotte Rancourt, a senior analyst with International Data Corp. (IDC), says that end users can expect to pay 30% to 50% less for SATA disk arrays versus high-end SCSI or Fibre Channel arrays, assuming equal capacity. For high-capacity configurations (greater than 2TB), the cost differential could be close to 50%, while smaller configurations (less than 4TB) would come closer to a 30% price difference.

At next month's NAB show, Ciprico will introduce a disk array based on Serial ATA that the company claims provides Fibre Channel performance at half the price. (Ciprico also sells

Fibre Channel disk arrays.) The company's FibreSTORE 2212A disk array is priced at \$17,100 for a 1.5TB entry-level version. Target applications for Ciprico's SATA disk systems include non-linear editing and film scanning.

Also at NAB, Ciprico will showcase its DiMeda (Digital Media Appliance) 1700, a network-attached storage (NAS) file server based on Serial ATA drives. (The DiMeda 1700 with SATA drives was introduced at last September's International Broadcasters Association show, where it won a STAR “best-of-show” award.)

The surge in popularity of SATA drives in non-desktop systems is expected to be almost as stunning. Reinsel predicts that shipments of “enterprise-class” SATA drives will go from half a million units last year to two million units this year and 3.6 million units in 2005 (see charts).

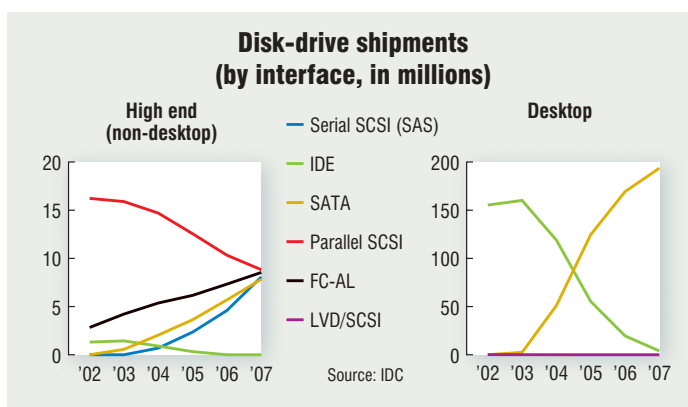
In addition to low cost, Serial ATA disk drives hold a lot of capacity per drive—250GB in most cases—relative to SCSI and Fibre Channel disk drives.

## PERFORMANCE FACTORS

On paper, Serial ATA is not as fast as Fibre Channel or high-end SCSI disk systems, but its performance may be sufficient in many media environments. SATA is rated at 150MBps, compared to 320MBps for Ultra320 SCSI and 200MBps for Fibre Channel. However, those are only theoretical maximum (“burst”) throughput ratings. Your mileage will vary.

Actual performance depends largely on the type of I/O (input/output) streams you're transferring. For random I/O with relatively small blocks of data, Serial ATA disk arrays will almost always be slower than Fibre Channel or high-end SCSI arrays. However, if you're transferring sequential I/Os in large blocks (for example, streaming video), the performance of a SATA disk subsystem may be almost as fast as a Fibre Channel or SCSI disk array. Fortunately, the latter scenario describes the I/O activity at most studios.


“Media environments are often characterized by very large, sequential reads and writes, and in those environments SATA delivers almost as much performance as Fibre Channel,” says Jim



(For a review of other recent SATA-related product announcements, see the New Products section in this issue, p. 49.)

As evidence of how popular SATA disk drives will be, IDC analyst Dave Reinsel predicts that SATA will be the dominant disk drive interface for desktop platforms within two years, with shipments skyrocketing from 2.5 million units last year to more than 51 million units this year, topping 124 million units in 2005—a compound annual growth rate of 406%.

**Q:** Over the next six months, the majority of your disk-array products will be based on which disk-based interface?



Source: InfoStor QuickVote reader survey

Farney, senior marketing manager at SGI (and the owner of a post-production facility—Pacific Video Resources, in San Francisco—for more than 20 years before he joined SGI).

“There’s a huge demand for Serial ATA in graphics-intensive environments, driven mainly by SATA’s low cost,” says Farney. (SGI started shipping SATA disk arrays, which are based on systems from LSI Logic Storage Systems, in the fourth quarter of last year. Early users of SGI’s SATA disk arrays include EFILM and Pacific Title & Art Studio, both of which are using the arrays for nearline archiving.)

Although Farney says that Serial ATA RAID arrays can fulfill the bandwidth requirements of most studio applications, if you need the highest bandwidth and/or you have a lot (20 or more) of users simultaneously accessing the same storage subsystem, you may need to go with Fibre Channel RAID arrays.

Of course, performance is very application-dependent and varies according to factors such as the size of the blocks being transferred, sequential vs. random access, and standard-definition vs. high-definition video, etc. But vendors that sell storage systems with a variety of disk interfaces (e.g., SATA, Fibre Channel, and SCSI) are surprisingly consistent in their performance evaluations.

“Our 12-drive SATA array is approximately performance-equivalent to our 10-drive Fibre Channel array—at about half the price,” says Gerry Johnson, senior product manager at Ciprico. (Ciprico discontinued its SCSI-based disk arrays about four years ago.)

#### RELIABILITY ISSUES

Because of its lineage as an interface for desktop platforms, SATA disk drives are often characterized as being less reliable than Fibre Channel or SCSI drives. However, vendors that have experience with all of the major disk drive interfaces say that, in environments that do not re-

quire constant storage access, there may be little difference in reliability among the various interfaces. “Even with our Parallel ATA [the predecessor to Serial ATA, also known as IDE] disk subsystems the drive failure rate was about the same as SCSI and Fibre Channel,” says Ciprico’s Johnson. “Serial ATA blurs the lines between desktop-quality drives and so-called ‘enterprise-class’ drives.”

While Serial ATA is creating a lot of buzz in graphics-intensive environments, some storage vendors—and end users—are sticking with other disk interfaces, at least for now.

For example, Huge Systems, a storage vendor that specializes in video and rich media markets, will showcase its line of SCSI disk arrays, including a new system based on the Ultra320 SCSI disk interface, at next month’s NAB show.

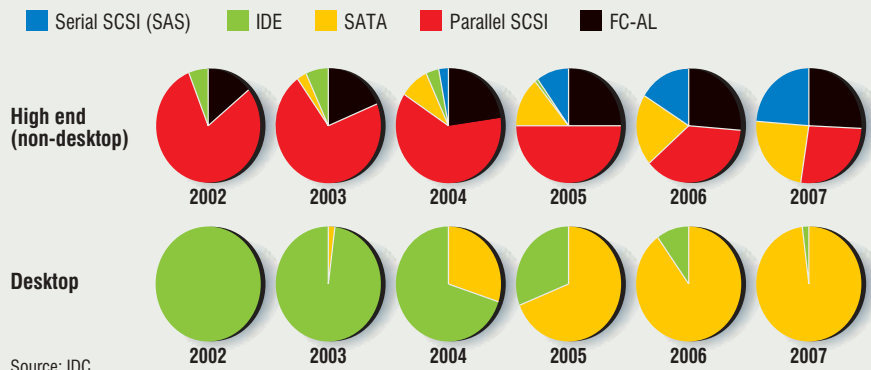
Mike Anderson, Huge Systems’ CTO, points out that Ultra320 SCSI may be the best fit for studios working with high-definition video. Anderson says that single-channel Fibre Channel disk arrays typically can’t handle high-definition video with a single link.

rays, Huge Systems plans to introduce a disk array based on Serial ATA drives in the third quarter. “The next generation of SATA drives will have better reliability, which will make them more appealing for 24x7 operation,” says Anderson.

Although some vendors pitch Serial ATA disk systems as primary (online) storage devices, SATA is also a good fit as a secondary storage device. In these scenarios, the SATA disk system sits between primary storage arrays (e.g., high-performance Fibre Channel arrays) and tape systems that are used for archiving content. Essentially, the secondary disk array allows users to “park” content on an inexpensive storage medium. The SATA disk subsystem is less expensive than the primary disk systems and is much faster than tape libraries for content retrieval.

“A film with a lot of special effects can be 30TB, or much more, and it doesn’t make economic sense to put that much capacity on expensive Fibre Channel arrays,” says SGI’s Farney. He adds that some large studios are starting

#### Disk-drive interface market shares



Source: IDC

Although Fibre Channel is rated at a theoretical maximum transfer rate of 200MBps, users typically realize only about 140MBps to 150MBps, which isn’t fast enough for high-definition video applications, according to Anderson. (You can add channels to the disk subsystem to boost throughput rates, but that increases the price of the disk array significantly.)

Anderson says that Huge Systems’ five-drive Ultra320 SCSI disk array can handle 220MBps, which is more than enough performance for high-definition video streams. The system can record and play back high-quality video (including 1080i HD at 10 bits) or can handle up to six uncompressed standard-definition video streams.

Despite its current emphasis on SCSI disk ar-

to implement “tiered” storage architectures where, once a film is completed, the content is moved to Serial ATA arrays for subsequent access and then to tape for long-term archiving.

In a related trend, some studios are using Serial ATA disk arrays to replace slow-access tape libraries as backup devices. SATA disk arrays are almost as inexpensive as tape libraries, but they provide much faster access times.

“There’s a big trend toward ATA and Serial ATA in studio environments, but mainly as secondary storage,” says Jay Krone, director of CLARiiON platforms marketing at EMC. “Digital content professionals can keep massive amounts of content ‘nearline’ or archived. ATA is also good in applications such as streaming video and video-on-demand.” □

# Studios move to network storage

*In media environments, there's a clear trend away from direct-attached storage to network storage, which provides a number of advantages for collaborative projects.*

By Dave Simpson

When it comes to the storage infrastructure underpinning your digital content operations, you have three choices: direct-attached storage (DAS), network-attached storage (NAS), or storage area networks (SANs). And in studio environments, the trend is clearly away from DAS toward networked storage architectures such as NAS or SAN (see Figure 1).

Jim Farney, senior marketing manager for media industries at SGI (and the owner of a post-production facility—Pacific Video Resources, in San Francisco—for more than 20 years before he joined SGI) says that most large studios

have already begun the transition toward networked storage, while smaller studios are still primarily using DAS. But he says that all sizes of studios can benefit from networked storage.

*“There’s a growing realization that professionals in collaborative environments have to move to some form of networked storage.”*

—Jim Farney

time-consuming task. There’s a growing realization that professionals in collaborative environments have to move to some form of networked storage, whether it’s NAS, SAN, or a combination of them.”

In addition to the content-sharing advantages of networked storage, studios may need the increased bandwidth of NAS (which attaches to Gigabit—1Gbps—Ethernet networks) or SANs (which are based on high-

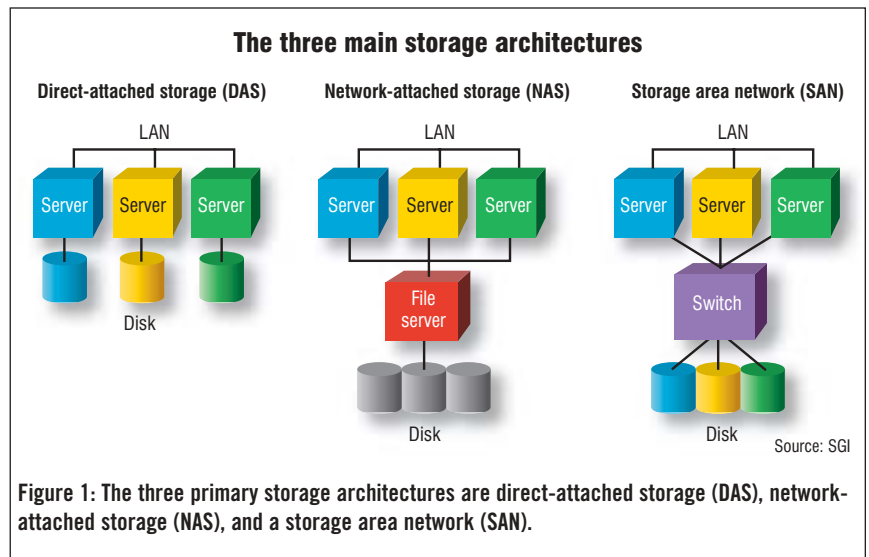
speed 2Gbps Fibre Channel networks). “In the broadcast media space there’s a big trend toward networked storage because media facilities are moving huge amounts of content, which requires very high bandwidth,” says Jay Krone, director of CLARiiON platforms marketing at EMC. “A 90-minute feature film in uncompressed HDTV format could require more than a terabyte of storage capacity.” (EMC’s CLARiiON storage systems were used by editors on a variety of films, including *The Perfect Storm*, *Gladiator*, *Mission: Impossible 2*, *American Beauty*, and *X-Men*.)

The trend toward networked storage is also being fueled by analog-to-digital conversion, which drives the need for much higher bandwidth and storage capacities.

## NETWORK-ATTACHED STORAGE

For digital content professionals, the primary advantages of NAS are low cost (versus SANs) and the ability to share files (versus DAS). Virtually all NAS servers support CIFS (the network protocol for Windows platforms) and NFS (for Unix platforms), and some NAS servers also support protocols for file sharing in Apple and other environments.

Although DAS is fine for professionals who do not need collaborative



**Figure 1: The three primary storage architectures are direct-attached storage (DAS), network-attached storage (NAS), and a storage area network (SAN).**

file sharing, or who spend most of their time working on still creation, for instance, NAS provides advantages in collaborative environments where high-speed file sharing is necessary.

Whether to migrate to NAS or SAN depends largely on your budget and bandwidth requirements. "NAS works great if you're doing, for example, standard-definition digital video editing with four or five video streams at up to 50MBps per client," says Gerry Johnson, senior product manager at Ciprico, a storage vendor that specializes in digital media applications. "But NAS may not provide enough performance in film editing environments that require throughput of, say, 200MBps and 2K-by-1K film resolution playback in uncompressed format."

NAS file servers are typically attached to Ethernet networks. Gigabit Ethernet is rated at a maximum of 125MBps, although in reality users may

achieve only 60MBps to 70MBps due to network overhead. According to a variety of storage vendors that focus on media applications, Gigabit Ethernet is a good choice when many clients must access files but none have sustained requirements for large

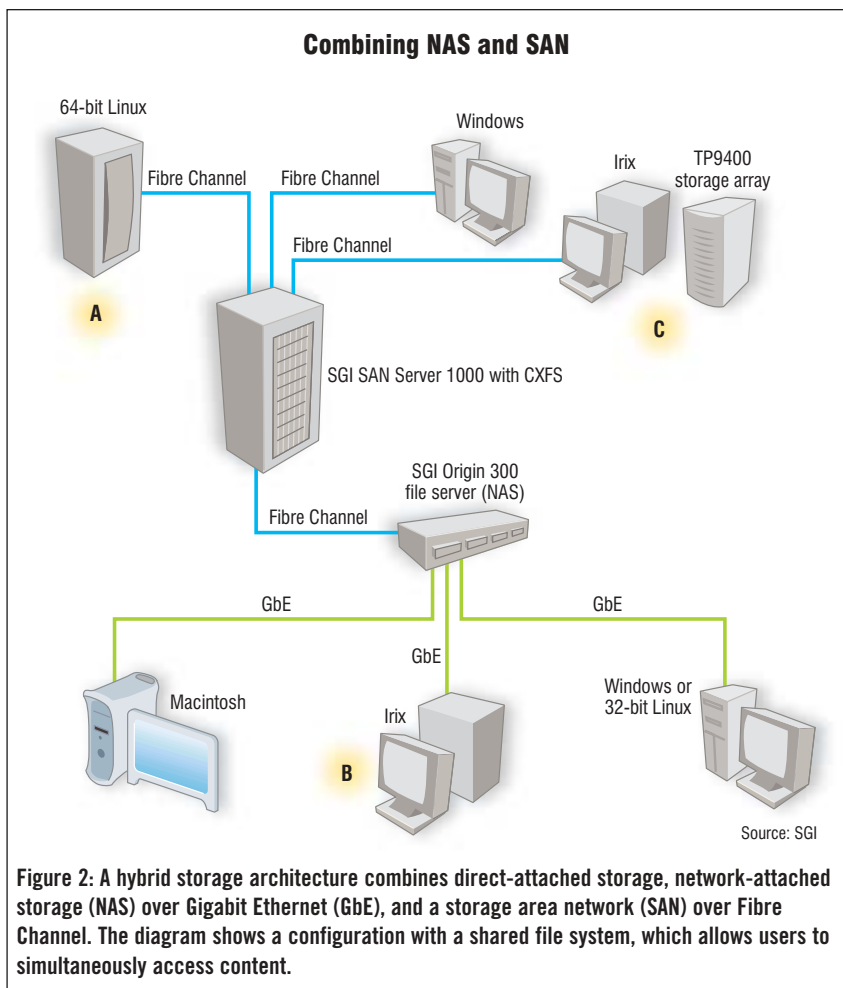
amounts of data, as in a render farm.

Gigabit Ethernet networks have boosted the throughput you can get with NAS servers and have enabled high-speed file sharing in media applications, but the real breakthrough will come with 10Gbps Ethernet, which is just starting to be used in IT environments. "Although expensive, 10Gbps Ethernet will enable more than adequate performance in all video environments, including high-definition video," says Ciprico's Johnson.

Another way to boost performance of NAS systems is to add TCP/IP offload engine (TOE) cards, which are available from vendors such as Alacritech and Intel. TOEs accelerate NAS performance by offloading TCP/IP protocol processing from the host processor.

Is NAS ready for big-budget film and TV production? To take just a couple of examples, Weta Digital Effects (in New Zealand) used NAS file servers from Network Appliance as primary storage for the special effects in New Line Cinema's *Lord of the Rings* trilogy.

Specifically, Weta used Network Appliance's F840 NAS systems, which Network Appliance refers to as "filers," for special effects and other content creation jobs involving terabytes of storage capacity, high-speed image manipulation (with single images that often exceeded 12MB), and more than 150 graph-



**Figure 2: A hybrid storage architecture combines direct-attached storage, network-attached storage (NAS) over Gigabit Ethernet (GbE), and a storage area network (SAN) over Fibre Channel. The diagram shows a configuration with a shared file system, which allows users to simultaneously access content.**



Network-attached storage (NAS) file servers from Network Appliance were used by Weta Digital Effects' graphics artists to create many of the special effects in the *Lord of the Rings* trilogy. *The Lord of the Rings: The Return of the King*. Photo: New Line Productions copyright 2003.

## Tippett banks on a SAN

Tippett Studio, a visual effects facility in northern California, faced huge amounts of data as it was working on films such as *The Matrix: Revolutions*, *Stepford Wives*, *Hellboy*, and *Starship Troopers 2*. To cope with the rapidly increasing storage capacities required for those projects, the studio installed a storage area network (SAN) in mid-2003.

"The main reason we went to a SAN was so that we could back up all of our data without impinging on server bandwidth," says Christian Rice, director of technology at Tippett. "Our disk capacity is doubling every year or so, and sometimes we generate more than 200GB of new data every day."

At the heart of the SAN are two 16-port, 2Gbps Fibre Channel switches from Brocade, which attach to the studio's servers and storage arrays (most of which are from SGI), and backup systems. The SAN has hundreds of client workstations (mostly Linux and Windows), and a total storage capacity of 16TB, which is expected to increase to 19TB in a month or two.

Rice says that in addition to enabling better backup operations (without negatively impacting overall performance), the SAN provides centralized storage and content management, as well as data sharing via SGI's CXFS shared file system. "All of our systems can access data simultaneously," he says. The SAN also allowed Tippett to balance workloads across the storage network to optimize performance and "to export data to clients much better and more dynamically," according to Rice.

Prior to migrating to a SAN architecture, Tippett's storage infrastructure was based primarily on direct-attached storage with some network-attached storage (NAS) systems.

ics artists. (Network Appliance NAS filers are also used to store the content at [www.lord-of-the-rings.com](http://www.lord-of-the-rings.com).)

Reel FX Creative Studios, a 3D animation and visual effects studio in Dallas, relied primarily on NAS file servers in a variety of film and TV projects, including the NBC movie *Saving Jessica Lynch*, heavily animated TV commercials for Hasbro, and a direct-to-DVD G.I. Joe film. Reel FX uses SGI's InfiniteStorage NAS 2000 systems for animation, effects, video editing, and post-production work in an environment with a variety of platforms, including Windows, Macintosh, Linux, and Irix—all of which can access the same

files simultaneously without file copying.

In the end, however, it may be the low cost of NAS that is most appealing for studios with limited budgets. For example, Huge Systems, a storage vendor that specializes in the video and rich media markets, sells a 1TB NAS server for \$10,600; a 3TB version is priced at \$23,900. The company targets digital video and standard-definition applications with its NAS servers, according to Mike Anderson, chief technology officer at the company. (Huge Systems will be demonstrating its NAS systems and disk arrays at next month's NAB show.)

### STORAGE AREA NETWORKS

If NAS doesn't provide enough performance, you may need to consider a SAN, most of which are based on Fibre Channel technology. "If you have hundreds of seats in a collaborative environment with high-definition editing, then a SAN is the only way to go," says Ciprico's Johnson.

For studio environments, the primary advantages of SANs are consolidated storage, centralized management, and a dedicated high-speed network. 1Gbps Fibre Channel is rated at a maximum transfer rate of 100MBps, although actual speeds may be more in the 75MBps to 80MBps range. 2Gbps Fibre Channel is rated at a maximum 200MBps, although speeds of 150MBps to 160MBps are more likely.

The downside to SANs is that they're expensive. The cost of the SAN hardware and software could easily eclipse the total cost of your workstations, graphics cards, and software. Pricing varies widely, but Fibre Channel host bus adapters (which are required for all servers and workstations attached to the SAN) can cost around \$1,000 per platform. And Fibre Channel switches are considerably more expensive than Ethernet switches—more than \$600 per port. In addition, Fibre Channel SANs require new cabling, specialized (and expensive) expertise and, possibly, expensive SAN management software.

However, SANs do provide the ultimate in storage performance, and when coupled with shared file systems, they provide file sharing (including concurrent file access) for collaborative work environments.

Shared file systems are available from vendors such as ADIC (StorNext File System), IBM (To-

*"NAS-SAN convergence will be a big trend in studio environments with a mix of bandwidth requirements."*

*—Jim Farney*

talStorage SAN File System), and SGI (InfiniteStorage CXFS). In addition, bundled Fibre Channel SANs with shared file systems are available from a few other vendors that specialize in the video/graphics markets, such as Avid and Pinnacle Systems.

Some shared file systems provide content sharing among diverse platforms, such as Windows, Unix, and Macintosh—a platform mix that is common in many studios. For example, SGI's CXFS file system supports Windows, AIX, Linux, Irix, Solaris, and Macintosh platforms.

The NAS-versus-SAN decision also relates to the type of content you're storing. "NAS is good enough for 8-bit standard-definition material, but if you need to share high-definition content then a SAN may be the only way to go," says Huge Systems' Anderson.

Although SANs provide high bandwidth and content sharing, some studios install a SAN for simpler reasons—at least initially. For example, Tippett Studio deployed a SAN primarily to solve backup problems, according to Christian Rice, director of technology at Tippett (see sidebar, "Tippett banks on a SAN").

### MIX 'N MATCH

The various types of storage architectures (DAS, NAS, and SAN) are not mutually exclusive. For example, you can mix NAS and SAN according to bandwidth and budget requirements (see Figure 2). "Professionals who need really high bandwidth, such as render farms, can use a Fibre Channel SAN while everybody else uses less-expensive NAS on existing Ethernet networks," says SGI's Farney. "NAS-SAN convergence will be a big trend in studio environments with a mix of bandwidth requirements. You get the file-sharing benefits of NAS with the higher bandwidth of a SAN."

A variety of vendors sell bundled solutions that combine elements of NAS and SAN. EMC, for example, has software that enables users to combine its Celerra NAS servers with its SAN-attached disk arrays. "There's a strong trend in the media industry toward combining NAS and SAN because it gives you the ease-of-use of NAS with the performance of a SAN," says EMC's Krone. □