

Optimizing Storage For The Media Creation Lifecycle

By Anand Singh Bisen
Principal Solutions Engineer
DataDirect Networks

Given consumers' current insatiable appetite for high-quality images, where facility operations managers no longer bat an eye at having to manage and store multiple Petabytes of data, the need for an elegant, eco-friendly and highly reliable shared storage architecture to handle this ever-increasing volume has never been greater. Deployed correctly, collaborative workgroups can produce more content, faster and with greater efficiency.

Yet it's difficult to find a single solution that fits every application. Indeed, every workflow brings with it a different set of challenges that must be overcome in order to make the most of the storage capacity available. To be valuable and streamline production and/or distribution operations, any storage system design must be fully optimized in order to accommodate and sustain a high level of performance, capacity and overall digital asset management.

With their need for fast-paced content creation, broadcasters and professional production companies must employ multi-tiered, block-level storage systems that deliver large amounts of bandwidth and unlimited, concurrent access to central storage repositories. The solution must also be easily scalable and accommodate all of the various types of existing and emerging storage technologies to extract the most value out of an organization's digital video and audio assets.

This storage architecture must also support multiple distribution platforms, like mobile video devices and the Internet, in addition to traditional TV outlets. Those engineers building out new storage system designs must first understand their individual workflow and then design a system that supports that workflow in the most cost-effective way.

Identifying Storage Requirements

For example, within a digital intermediate (DI) workflow for motion picture production, the data rates for a single real-time stream is close to 300 megabytes (MB) for every second of footage. For a 120-minute feature, it is not unusual that the production will shoot over 800 minutes of feature.

In terms of storing and processing those images, they are faced with the challenge of capacity and performance. In order to share, view and process (color grading, compositing, etc.) this high-resolution footage, they require a storage subsystem powerful enough to sustain multiple of these streams in real-time concurrently.

Therefore, DI professionals need storage systems that can deliver multiple gigabytes of data per second of sustained throughput and allows them to plug into a central storage system and work on the same file simultaneously, without bottlenecks.

Broadcasters, on the other hand, are not capturing content at such high quality. They typically work with compressed formats in order to conserve bandwidth. There are many channels to support, so the need for concurrent access to the content is much greater. Yet, the bandwidth requirements for editors and producers are less than with DI workflows. Instead of a small amount of large data

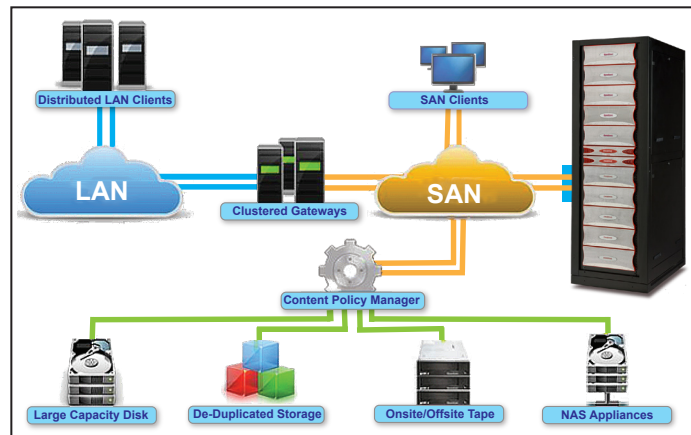
streams, broadcasters need hundreds of streams at 40-50 MBps. So, the biggest challenge is to make available a high level of overall concurrency to accommodate dozens of networked workstations.

Another critical aspect of the broadcast business is play-to-air support. During large live, multi-venue events (like the Olympic Games), where you have a need for a lot of remotely located editors to work collaboratively, there is a real-time requirement that is critical to success. You need to add commercials and animation to the

live stream. In this situation, you need to make sure that you are sustaining a certain level of performance when a bunch of editors are concurrently ingesting content in real time, adding new content (commercials and IDs) and pushing it out through the other end; while also interfacing with a content management system.

Shared Asset Efficiency

Most of the current storage deployments for broadcast applications are configured as storage area networks (SAN) because the bandwidth requirements on a per client basis are very high but the latency is low. In many cases the SAN is complemented by a Network-Attached Storage



DDN technology provides flexibility of connection to give the highest bandwidth applications direct connection to disk, while also offering scalable Ethernet-based access for lower-resolution streaming, connecting virtually every operating system into a single scalable volume and enabling data lifecycle management.

(NAS) array that supports collaborative workgroups of editors and others involved in the content creation process. That's because these creative people need to perform a large number of file transfers.

Fox Cable Network's SPEED Channel has deployed such a SAN to support a complex file-based workflow at its 55,000 square foot production facility in Charlotte, N.C. The system, based on an xSTREAMScaler File Storage Platform from Data Direct Networks (DDN) and a shared-storage editing environment from Dalet Digital Media Systems, facilitates the seamless transfer of SD and HD content throughout the plant. Footage is ingested directly into the SAN, where low-resolution proxies are created for easy editing, and file sharing to and from dozens of video servers. The innovative architecture now allows 75 concurrent Dalet users to go perform everyday tasks, from logging content to rundown preparation.

The xSTREAMScaler system provides SPEED's production team with real-time access to digitized content while simultaneously allowing all applications and devices to work seamlessly without worrying that one application may cause disruptions such as dropped frames or delays.

Self-Healing Assurance

DDN platforms support an aggregated level of streaming content, which reduces the total cost of ownership, and also maintains real-time performance while the system is recovering from various errors or failing disk drives.

This uninterrupted system performance is made possible by a series of proven technologies that provide smooth scalability and reliable data integrity without compromise. This includes performing self-healing tasks in the background should a drive go bad, or even lose up to 20 percent of the disks and still provide uninterrupted performance without any degradation in performance for the networked workgroup sharing that storage.

Data protection is also critical. Using SATA disks, for example, given the large number of read operations, the data may be correct on the underlying storage, but by the time it comes off the disk, there can be in-flight bit errors. Because the data is stored over an extended period of time and users will be reading the same data and re-provisioning it hundreds of times, as a result the quality of data over a period starts decaying. DDN has developed a technology called SATAAssure where the system verifies and corrects the content using the parity information on every read operation to remedy this and maintain peak performance at all times.

With all of this data that is being generated at every phase of the workflow, a massive archive not only becomes crucial in terms of the viability and security of the data, but also in terms of a central repository where you can save

all of this data reliably and in very cost-effective manner. However, the more disks you have in your system, the chances for disk failure increase exponentially.

With the S2A9900 and SFA10K systems, not only is each disk RAID protected, but the back end channels and entire enclosures are RAID protected as well. This allows the system to not only withstand a single or multiple drive failure, but the entire system (enclosure, back-end channels) in a redundancy group can go down and the system continues to provide uninterrupted access to the storage without any performance hit. Sustained performance during a catastrophic situation is more important than reliability of a single drive. Leveraging a technology called DirectRAID, all DDN storage solutions instinctively rebuild corrupted data and flush the repaired data back to the disks to self-heal the problem automatically.

Because DDN systems like the S2A6620 and S2A9900 can fit more than a petabyte in a single rack, they not only provide the user with very fast access to the data but also do so with very low power consumption in a dense package. This helps to save on electricity and real estate, and makes for a very environmentally friendly solution. This unique density and power savings brings a fast return on investment and is made possible by a disk power down solution, called Dynamic Massive Array of Idle Disks (D-MAID), which is available on all of the company's data infrastructure platforms. D-MAID allows customers to dynamically spin down storage LUNs (logical unit numbers) and enter a power savings mode based on user-specified parameters. This greatly reduces power consumption and operational costs.

Success Means More Than Just Storage

Because broadcasters have to deal with so many issues when it comes to data storage and availability, clients today are not buying storage; they are trying to solve a unique problem that includes storage as one element of streamlining a particular workflow. That's why DDN considers itself to be more than just a storage vendor. We add value to facilitate seamless interoperability with different applications from third-party companies—in the form of database management and file distribution and sharing technologies—to deliver the performance and problem-solving solutions our customers need.

DDN understands that there can be no bottlenecks anywhere along the content creation chain. We know how to optimize the different types of underlying storage solutions to get the most value for our customers and we work with our customers to customize systems for different applications. This system engineering expertise is not a separate service, but part of our commitment to assure that our customers are successful. The goal is to make the storage as useful to the production team as any of the individual applications tied to that storage.

[Click here to see more of DDN's Media applications.](#)

DataDirect[™]
NETWORKS

ddn.com
1.800.TERABYTE

9351 Deering Avenue
Chatsworth, CA 91311
USA