



Entertainment and Media's Post-production Storage Crisis

Abstract

Collaboration is an important requirement for post-production editing, compositing, animation, special effects, color grading, mastering, and archiving. SAN storage does not natively provide that collaborative workload sharing. Workload sharing can only be run on a small high performance SAN environment when there is a shared SAN file system. When those shared SAN file systems are scaled to a mid-sized or larger SAN, system performance rapidly degenerates. To effectively scale this type of SAN system means providing dedicated systems to small groups.

This motivates the vast majority of production organizations to quickly push real-time content from SAN storage onto much lower cost high performance NAS for their post-production work. The good news is that NAS greatly reduces post-production collaborative storage costs. However, most NAS systems cannot meet the demands of post-production. They can't effectively scale performance or capacity and end up causing serious performance and operational bottlenecks. The result is a post-production storage crisis.



Introduction

In the world of “real-time” uncompressed video and digital film, high performance SAN storage (4 or 8Gbps Fibre Channel and 10 or 20Gbps Infiniband) is an absolute requirement. Yet, one very important consideration to high performance SAN storage and infrastructure is the extremely high costs. Another is its limited ability to scale concurrent shared workload access per system.

Collaboration is an important requirement for post-production editing, compositing, animation, special effects, color grading, mastering, and archiving. SAN storage does not natively provide that collaborative workload sharing. Workload sharing can only be run on a small high performance SAN environment when there is a shared SAN file system. Shared SAN file systems are typically available from Apple, Avid, Bright Systems, and SGI. When those shared SAN file systems are scaled to a mid-sized or larger SAN, system performance rapidly degenerates. The cause is primarily the increased load on the metadata controller providing too few resources to guarantee the data stream for video/film real-time streaming. To effectively scale this type of SAN system means providing dedicated systems to small groups.

This motivates the vast majority of production organizations to quickly push real-time content from SAN storage onto much lower cost high performance NAS (file storage that inherently provides workload sharing) for their post-production work. The good news is that NAS greatly reduces post-production collaborative storage costs. It allows the post-production storage to represent about 80% of the capacity for approximately 50% of the total storage budget. Unfortunately, most NAS systems cannot meet the demands of post-production. They can't effectively scale performance or capacity and end up causing serious performance and operational bottlenecks. The result is a post-production storage crisis.

Why The Use of Most NAS Creates a Post-production Storage Crisis

Performance

Traditional NAS, and even much of the newer clustered NAS, just is not designed to handle the performance requirements of post-production. They are typically incapable of concurrently providing both high IOPS and high throughput. And even when the NAS system is tuned for one metric or the other, it still comes up short. For example, most NAS systems tend to have extremely long rendering times with read access to the digital assets that is much too slow and write out that is too slow after the rendering job has completed, forcing the artistic talent to wait. Often, the systems cannot provide backup during the day, further encroaching on rendering times. There is also frequently a severely constrained ability to combine rendering and workstation workflows in the same time frames.

These performance constraints can lead to application failures or “hangs” as a result of I/O timeouts.

Performance Scaling

Scaling performance is another critical NAS system issue. As clients increase, the NAS systems have acute performance limitations. Swelling number of clients causes desktop response times to crater (becoming much slower) upon opening or saving a project. These performance issues only get worse when meeting the requirements of HD (6MB), 2K (12MB) and 4k (49MB) files.

One of the more troublesome performance scaling issues is when the NAS system runs at its maximum performance or maximum files, it can and often does become non-responsive. Far too often, IT organizations have to deal with restarting the system and recovering data. These crashes are non-trivial events that cause significant user disruptions. When data is lost, it has to be rebuilt or recovered. Getting that NAS system back up, restarted and running, can be a very lengthy painstaking process. The primary work-around is to implement more NAS systems. However, that is not a very good work-around. It doesn't solve the problem, but merely defers it while creating other problems of its own.

Thorny Consequences

There are two significant consequences to these performance issues. The first is the inescapable decline in expensive talent productivity. In the example of slower desktop times, the talent is distracted from their creative workflow while their scenes load or save out. They tend to end up grabbing coffee and/or socializing with and generally distracting the other talent. When there are NAS system performance issues, the talent simply cannot perform their work in a timely effective manner.

The second consequence is more worrisome, with far-reaching implications. NAS systems' inability to effectively scale within a system results in increased numbers of NAS systems. That's more NAS systems to implement, operate, and manage, while providing rack space, power, and cooling. Unfortunately, that's the least of the problems.

NAS management does not increase in linear proportion to the increase in systems. It increases exponentially. Increased numbers of NAS systems require redundant and pedantic ongoing basic management. Then there is the constant hands-on data migration between systems plus continuous recalculated load balancing, if there is to be any semblance of efficient asset utilization. This nightmare gets exponentially worse as the number of systems continue to grow.

That's when it gets even worse. More systems equals more rack space, more power, more cooling. What happens when all that additional real estate, power, and cooling exceeds the space available? The entire organization, or at least the IT organization, must move to a new location. This is an incredibly time consuming and costly consequence all because the NAS system does not perform or scale effectively enough for post-production.

High-performance SAN storage with SAN file systems do not scale users or performance very well and are much too expensive for post-production work. The majority of NAS systems neither perform nor scale well enough to meet the requirements of post-production work. They end up requiring too many systems to get the job done radically increasing the actual total cost of ownership. This conundrum has led to Entertainment and Media's post-production storage crisis. There has to be a way to solve this conundrum. There is.

The BlueArc Titan Solution

Titan Performance

BlueArc designed the Titan series from the ground up to meet post-production requirements of scalable collaboration, performance, and capacity. It is architected for the best possible performance. Common performance draining storage functions such as CIFS and NFS have been implemented into the silicon. The BlueArc Titan series has time and again set the industry standard for ultra-low latency (measure of how quickly a server can respond to clients' requests) for server IO request handling. The lower the latency, the greater number of clients and IO that can be handled. The low latency capability concurrently allows for combined rendering and workstation workflow. Titan maintains this rapid response even under the highest demands as each part of the architecture handles the max load simultaneously with distributed memory, state engines and pipelines enabling massive parallelism. This parallelism permits the Titan to run effortlessly at its maximum performance limits, ongoing without any degradation or hiccups. Titan is a NAS system that requires no compromise between IOPS and throughput. It has unfailingly provided independently proven industry-leading performance in ongoing tests such as SPECsfs (<http://www.spec.org/sfs97r1/results/sfs97r1.html>).

Titan's unique high performance and low latency architecture allows it to provide astonishing improvements in rendering speeds, while noticeably reducing artist wait times from file transfers to downloading shadow masks, and increasing their productivity. And because Titan's high performance architecture removes data-flow constraints, it eliminates the performance bottleneck for even the largest most powerful render farms. The outcome is faster renders; fewer dropped



frames; and much less rework for the render wranglers. The Titan's incredibly fast digital asset access and write out after the rendering job has completed, allows artists to focus their attention on spectacular results and higher quality by providing much faster read access to the digital assets. In other words, productivity increases instead of decreases.

Titan Scalability

Titan provides incredible scalability that scales up – the scaling capacity within the NAS filer, and out – the clustering of more than two heads in a single image. Titan scales up by virtualizing backend RAID storage. All backend storage attached or added to each Titan is virtualized into a single or multiple pools. Since the Titan scales capacity by adding additional RAID controllers as well as hard disk drives, more processing power is available for the storage subsystem. More processing allows linear scaling of performance and capacity. Other NAS filers do not scale the backend RAID processing so that at a specific capacity point there are diminishing marginal returns where the performance actually declines as capacity increases.

Titan scales out with simple clustering of up to eight Titan nodes per system. In combination with the virtualized backend SAN storage, thin provisioning, and multi-head clustering, scaling a Titan “system” can be performed online while in use and non disruptively to the application.

The Titan's well thought out hardware architecture, hardware-based file system, extraordinarily high speed task processing, and minimal IO latency are all key factors that allow Titan to scale to unprecedented levels.

Each Titan node supports scalability capabilities of:

- Up to 12,000 Windows (CIFS) users
- Up to 60,000 UNIX (NFS v2 or NFS v3) users
- Up to 16 million files per directory
- Up to 8 billion total files
- Up to 200,000 IOPS (NFS v3)
- Up to 20Gbps aggregate throughput
- Up to 4 PB of total usable storage (per Titan or cluster)
- Up to 256 TB per file system

With up to eight Titans in a single clustered system image, Titan sets new standards in collaborative, performance, and capacity scalability. This allows Entertainment and Media organizations to architect a solution that today may only have only 200 rendering nodes and 40 workstations but will scale to 1,000 rendering nodes and 160 workstations within a year while being completely confident that Titan easily handles it.

Titan's ability to scale allows it to continually grow as more projects are awarded, more artistic talent is hired, and projects are developed in higher resolution.

More importantly, Titan's extensive scale up capability means a lot less rack space, less power, and less cooling is required. Titan's scale out (clustering) capability additionally means much simpler system management and no requirement for hands-on manual movement of data and load balancing.

Titan goes even further with an intuitive graphical user interface that makes Titan even easier to implement, operate, scale up, scale out, manage, adapt to the unexpected, and simply use. It makes both common and complicated operations simple often eliminating the requirement for a fulltime dedicated storage administrator.

Summary and Conclusion

The Entertainment and Media industry is going through rapid changes today. File sizes are growing exponentially with the implementation of higher resolutions (HD, 2K, 4K) requiring higher performance storage with highly scalable capacity. As desirable as high performance SAN storage and infrastructure is for real-time data capture, it is equally undesirable for post-production work. It is much too expensive and far too difficult to scale workload sharing.

The common answer for post-production work is NAS. It has built-in collaboration and workload sharing, costs considerably less than SAN storage, and is simple to implement, operate, manage, and use. Regrettably, most NAS systems simply cannot adapt the rapidly changing realities of the 21st century Entertainment and Media market. So they proliferate far more rapidly than expected, creating management, operational, real estate, power, and cooling problems of crisis proportions.

BlueArc identified these problems years before it became a crisis and developed the Titan NAS system to eliminate them. The Titan uniquely adjusts and adapts on the fly (online and application non-disruptively) to the ever-increasing performance and scalability requirements of post-production work. It eliminates the problems of system proliferation as well.

Every Entertainment and Media IT organization would be remiss if it did not consider using the BlueArc Titan series for post-production work. It is the surefire way to increase artistic talent productivity while eliminating a post-production storage crisis.



About BlueArc

BlueArc is a leading provider of high performance unified network storage systems to enterprise markets, as well as data intensive markets, such as electronic discovery, entertainment, federal government, higher education, Internet services, oil and gas and life sciences. Our products support both network attached storage, or NAS, and storage area network, or SAN, services on a converged network storage platform.

We enable companies to expand the ways they explore, discover, research, create, process and innovate in data-intensive environments. Our products replace complex and performance-limited products with high performance, scalable and easy to use systems capable of handling the most data intensive applications and environments. Further, we believe that our energy efficient design and our products' ability to consolidate legacy storage infrastructures, dramatically increases storage utilization rates and reduces our customers' total cost of ownership.



BlueArc Corporation
Corporate Headquarters
50 Rio Robles Drive
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202