Coughlin Associates
Data Storage Consulting

NVMe: Driving Changes in the Traditional M&E Storage Hierarchy

Tom Coughlin
Coughlin Associates
www.tomcoughlin.com

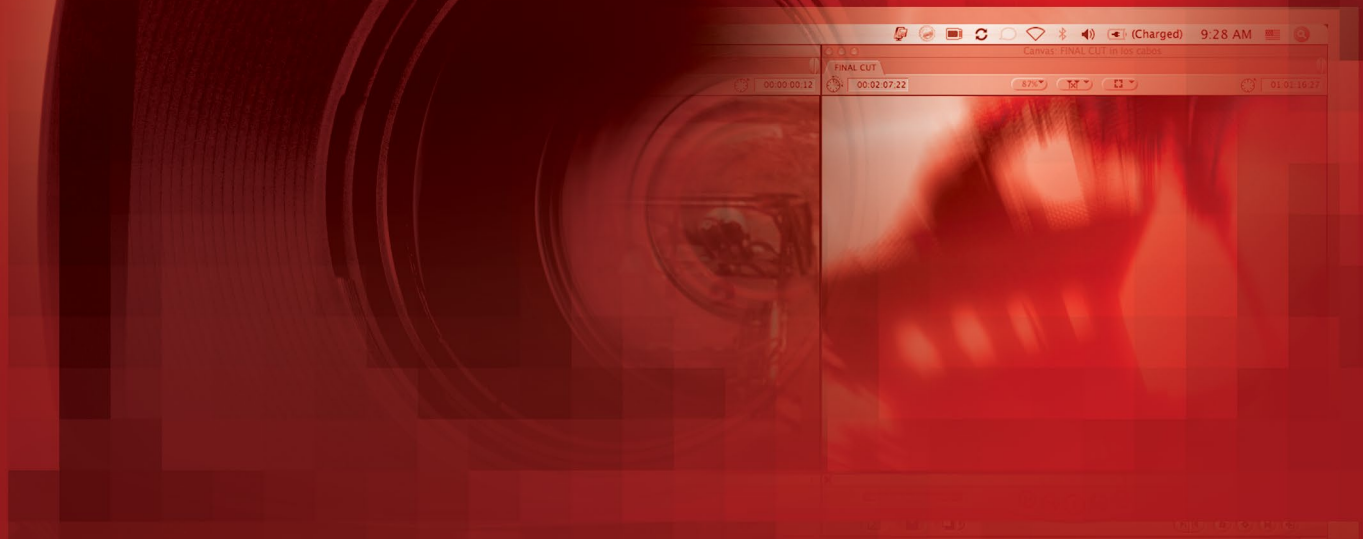

CONTENTS

MEDIA AND ENTERTAINMENT TRENDS

Media and entertainment storage and bandwidth requirements are increasing thanks to higher resolution content combined with higher frame rates, higher number of bits per pixel and more cameras per project. This richer content must be created cost effectively and with the most efficient use of physical and human resources possible.

In many ways the anticipated future is arriving faster than originally projected in the media industry. For example, NHK in Japan will be broadcasting the 2020 Olympics in 8K. Taking note of this trend, camera manufacturers are committing to bringing ultra high-resolution to the masses, with 8K cameras already available from Red Digital, Sony, Canon and several others. These cameras often capture content at multiple Gbps data rates and store more than a terabyte of digital content per hour. Besides impacting professional video camera storage and memory needs, these trends are also driving higher bandwidth and processing power for video rendering and special effects.

A 4K 360-degreee video project may use 4 to 16 cameras, each capturing a separate stream of 4K or even 8K content. After they are captured, multi-camera streams for 360-degree video must be stitched together, resulting in a much larger composite video file. Such an operation is heavily dependent on the underlying storage architecture, requiring considerable bandwidth and processing time.

The major post production software companies are also getting in on the action, supporting 8K workflows across numerous content applications. Even the upcoming Xbox Project Scarlett and PS5 game consoles will have 8K graphics support by 2020. And finally, 8K TVs are available from all the major consumer electronics companies.

In addition to games, higher resolution content is available on streaming and other distribution channels, as most media organizations understand this content is favored by consumers (assuming the quality of content delivery is acceptable). Taken together, this drastically impacts storage and bandwidth demands for content transcoding, IMF operations and content distribution.

For those in the media and entertainment industry tasked to work with and produce content at such high fidelity, it is increasingly apparent that new technologies must be leveraged to support these workflows. One such example is increasing efficiency by handling this content with advanced asset management tools – many of which now deploy AI algorithms to help users find and access media content. Another example is using advanced lossless and lossy compression technologies to reduce the size of delivered content. Finally, one of the most well-utilized and best ROI methods for dealing with these new requirements is to use the inherent capabilities of different types of modern storage media.

DIGITAL STORAGE OPTIONS FOR MEDIA POST-PRODUCTION

With the growth in both size and amount of digital content, total storage capacity for media post production is rapidly increasing. Figure 1 shows projections for the growth of local direct attached, local network and cloud storage for post-production applications out to 2024.

Figure 1. Growth in Digital Storage for Post-Production

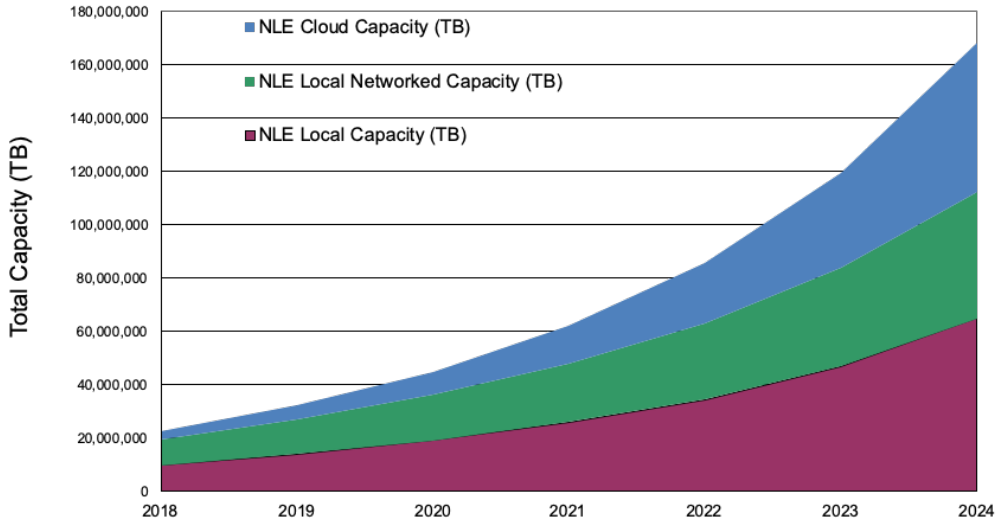

Hard disk drives (HDDs) are historically the most popular storage media in media post production facilities, sometimes with magnetic tape or storage in the cloud used to backup or archive media content. The cloud (tiered storage in data centers) is often used for storing and sharing content for collaborative media editing projects.

Although HDDs will remain an important storage media for parking content and storing content not frequently accessed, they have inherent performance limitations in the context of editing today's high capacity and high bandwidth 4K+ content. Thankfully, as manufacturing efficiencies have improved and competition has become more fierce, the economics of NAND flash memory has improved enough such that many organizations are now able to consider this storage media for a larger percentage of post-production workflows.

Like most developments in the technology industry, things rarely stay stagnant for long, and such is the case with flash memory devices. Although Serial AT Attached (SATA) and Serial Attached SCSI (SAS) SSDs are widely used, these older interfaces (originally created to support the needs of HDDs) can restrict the data rate and latency that SSDs are capable of achieving. This has led to the wide use of an interface that brings more of the internal performance of the SSD to the computers it is connected to. This new interface is called Nonvolatile Memory Express (NVMe).

NVMe SSDs

NVMe SSDs avoid many of the built-in delays in SATA and SAS interfaces. A key reason for this is that they are based upon the PCIe bus used in modern computer systems, which allows the devices to support many more commands per queue than either SATA or SAS. The PCIe bus roughly doubles the available per lane data rate about every three years, providing greater bandwidth for video workloads. PCIe 3.0 is being used widely today and provides 8 giga-transitions per second (GT/s) per lane while PCIe 4.0 is starting to appear in computer systems, providing 16 GT/s per lane. This roadmap for supporting technology means even higher data rates will be possible in the future.

NVMe SSDs come in several form factors and can use multiple PCIe lanes to provide higher performance. These form factors can look like 2.5-inch HDDs or be very compact to fit into a tablet-sized computer system. NVMe SSDs are used in mobile computers, tablets and enterprise storage systems. Figure 2 shows a 2.5-inch form factor NVMe SSD. This form factor is often used in enterprise storage systems.

Figure 2. 2.5-inch Form Factor NVMe SSD

From the perspective of performance, it's no surprise that NVMe SSDs are increasingly used as primary storage for many M&E applications. NVMe SSDs also provide lower latency to content than HDDs, which can significantly facilitate the creative process for media professionals.

However, performance is not the only reason organizations are making the switch to NVMe SSDs. With the lower price of SSD storage, their total cost of ownership has declined, making the lower capacity less of a trade-off than it was just a few years ago. Organizations are also benefiting from lower data center costs thanks to the reduced power and cooling requirements of SSDs, as well as the greater performance density they can offer. In some use cases, organizations are leveraging the capability of these devices to provide high-performance across an all-IP network, enabling them to switch from Fibre Channel to Ethernet-based infrastructure, further improving the overall economics of their storage environment.

Some M&E Uses for NVMe SSDs

One compelling use case for NVMe is to support video projects that involve high frame rate, high bit per pixel 8K video. Editing this content can require nearly 10GB/s streaming performance at sub 16ms latencies – performance which pushes traditional HDD storage to the limit. NVMe-based storage can provide low latencies and the bandwidth required to edit a raw 8K video project.

Another use case involves content created using multiple cameras capturing 4K+ video and then working with multiple streams of these high bandwidth assets. Because of its random IOPs performance, NVMe-based storage can provide the low latencies required to edit multiple Gbps data streams and to stitch together 360-degree multiple camera video.

Another applicable use case for NVMe is for visual effects and rendering tasks, especially given their intermittent data demand. All of the render engines in a render farm tend to get new data from storage at about the same time, making the peak bandwidth demand to efficiently service these render engines very high. NVMe-based storage can service the very high intermittent demand for 8K and higher rendered content more efficiently than an HDD-based storage system.

Best Practice Implementation of NVMe Storage Arrays

Because the raw storage capacity cost of HDDs is considerably less than SSDs, many people seek a less expensive storage option for their M&E workflows. However, there are good reasons to consider including NVMe SSD storage in your workflows, and one should consider more than the capital cost in deciding to use one type of storage over another. Following are several aspects of NVMe to consider that tend to break with conventional wisdom.

Cost: First, the cost for a given amount of SSD storage capacity has dropped considerably over the last year or so, making the purchase of some flash memory much more affordable. And because SSDs are built using a bunch of NAND flash chips on a circuit board, they can be made very small and dense, allowing more storage in a given data center footprint than HDD storage. In addition, NVMe SSDs can be built into storage systems using NVMe over fabric (NVMe-oF) where that fabric can be Ethernet. This enables NVMe storage systems to play an important role in IP-based workflows.

Lower Latency: HDD systems, even when they can provide the required data rate, have considerably higher latency than SSD systems. This can make a big difference for the efficiency of your creative professionals, allowing them to not only get more work done in less time, but also give them more creative freedom to produce even better content. Such improved throughput and content of creative staff can more than compensate for the additional cost of SSD storage.

Multi-Tier Benefits: Finally, SSDs can be a part of a multi-tiered storage system, allowing an optimal combination of less expensive HDD storage and higher performance NVMe SSD storage. Less active data stays on the HDDs and only active data is on the SSDs.

QUANTUM F-SERIES NVMe APPLIANCES

Quantum introduced its F-Series NVMe storage system to meet the needs of media professionals and power their modern post-production workflows. Figure 3 shows Quantum's F2000 dual-ported NVMe storage system.

Figure 3. Quantum's F-Series NVMe Storage Array

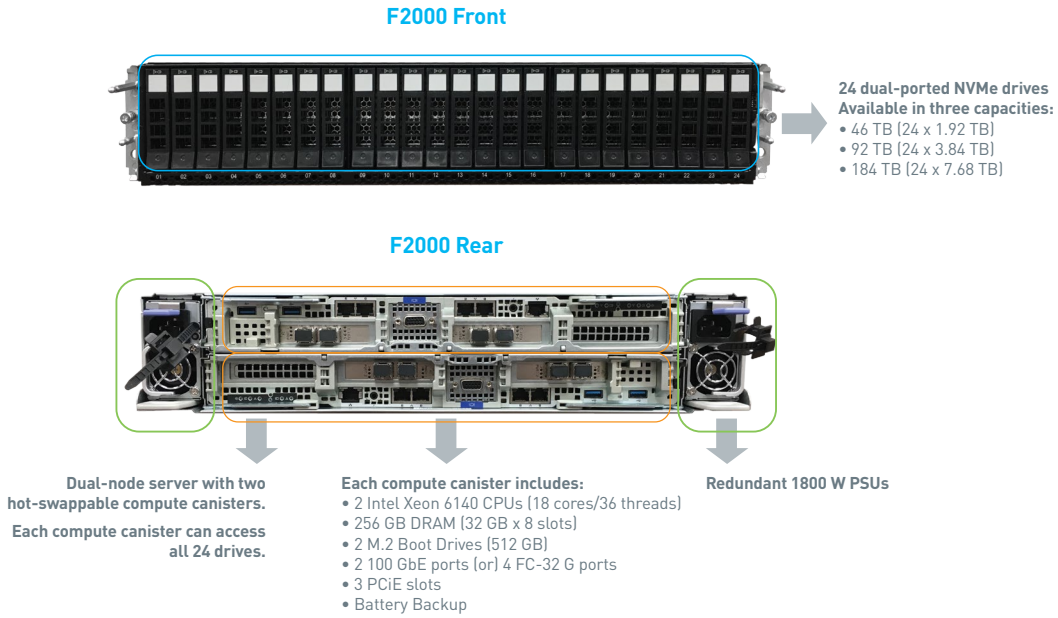

As shown above, storage capacities up to 184 TB are available with current SSD capacities. High uptime is ensured by dual ported SSDs, dual-node servers and redundant power supplies. The NVMe SSDs allow performance of about 1 million random reads per second with latencies of under 20 microseconds. Quantum found that their NVMe storage can deliver more than 10 times the read and write throughput performance with a single client compared with NFS and SMB attached clients; specifically, over 25GB/s in aggregate streaming bandwidth.

The F-Series was designed specifically for video and video-like data, and was made to handle the performance requirements of multiple stream 4K+, high frame rate data as well as other types of unstructured data. It uses 100% NVMe SSDs support a huge amount of parallel processing, as well as Remote Direct Memory Access (RDMA) networking technology to provide direct access between workstations and the NVMe storage devices.

These capabilities enable editors across several locations to work on multiple streams of 4K and even 8K video using one storage volume. The higher performance of NVMe SSDs avoids the overprovisioning often required with HDD-based storage systems.

The F-Series allows faster ingesting of source materials and higher performance editing, especially with multi-stream content. Transcoding is faster, as are animation rendering data calls. Playout of content and master generation are also faster. Smaller facilities might be able to store all their active content on NVMe SSDs, while larger facilities would likely use NVMe systems as part of a multi-tier storage environment.

NVMe CASE STUDIES

The Quantum F-Series NVMe storage array allows working with uncompressed 4K or 8K video, allowing end-to-end online workflows without having to transcode to compressed proxies and speeding the creation of deliverable masters. It also speeds post production using digital intermediaries that involve digital ingest, color grading and correction, dust busting and other steps.

By harnessing the power of the F-Series, a larger video production studio was able to get almost a 66% reduction in their storage arrays, as well as an 83% reduction in drive count and rack space. They also achieved a 285% increase in available bandwidth and a 35% reduction in power consumption.

Another large studio with a single F2000 was able to simultaneously read 18 streams of 2K standard (or anamorphic) resolution frames and 8 streams of 4K standard (or anamorphic) resolution frames. Figure 4 shows the observed performance as well as 8K write and read results.

Figure 4. Studio Single F2000 Observed Performance

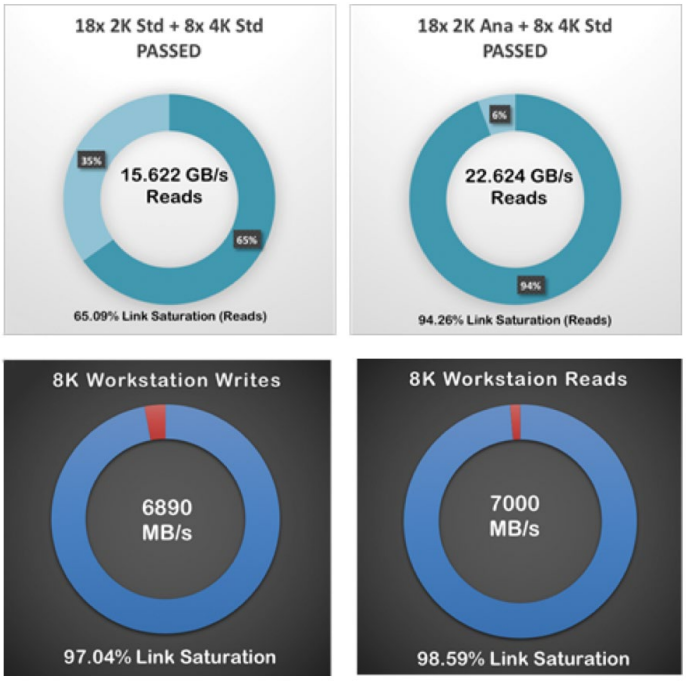

VFX is also faster with NVMe storage. A mid-sized VFX shop in Canada was using a QXS-656 hybrid SAS system combining SSDs and HDDs. They installed an F2000 array and were very impressed with the product. The customer said that their production test render would take almost 75 minutes to run on the QXS-656, but with the F2000 the same render job took only 7 minutes. The company now intends to run their entire production pipeline off the F2000.

Finally, modern workflow tools using machine learning and other types of AI, as well as IMF encoding on the fly and mixed reality workflows are faster using the F-Series, and will greatly benefit from the kind of performance NVMe can offer.

CONCLUSIONS

NVMe SSDs enable maximum flash memory performance and can be networked together using NVMe over fabric technology to create highly responsive storage systems. The growth in the size and volume of media content in modern media workflows - as well as higher streaming requirements for high frames - have made it harder for traditional storage solutions to provide the performance media professionals require. As a result, high performance low latency storage using NVMe technology is becoming an essential element for many post production, animation and VFX workflows.

NVMe storage increasingly serves as primary storage in many tiered media storage applications and, for some smaller facilities, SSD storage often comprises the bulk of their capacity. NVMe-based storage systems, such as the F-Series appliances from Quantum, allow modern studios to create high resolution content faster, render animation and VFX content quicker, enable new AI-based media tools, and free media creatives to spend their time mastering their content rather than fighting their equipment.

ABOUT THE AUTHOR

Tom Coughlin, President, Coughlin Associates is a digital storage analyst and business and technology consultant. He has over 37 years in the data storage industry with engineering and management positions at several companies. Coughlin Associates consults, publishes books and market and technology reports and puts on digital storage-oriented events. He is a regular storage and memory contributor for forbes.com and M&E organization websites. He is an IEEE Fellow, President of IEEE-USA and is active with SNIA and SMPTE. For more information on Tom Coughlin and his publications and activities go to www.tomcoughlin.com.

Coughlin Associates

Data Storage Consulting