



NVMe-oF: The New Frontier for On-set Production

Introduction

For Media & Entertainment, getting all the raw camera footage from on-set dailies into the pipeline to be transcoded, reviewed, archived and delivered as quickly as possible is an ongoing balancing act for solution stacks. The higher the frame size, frame rate and footage definition, the more data there is to shift.

Media & Entertainment has widely embraced commodity IT technologies, such as cloud and Ethernet, to do these processes. However, Storage Area Network or SAN is still trusted to deliver real-time playback across many disciplines, workflows and pipelines. It is a widely held assumption that as formats and data-rates become ever higher, a SAN is needed.

This whitepaper will explore the 'why' behind this trust, and how the introduction of new technologies, such as NVMe over Fabrics or NVMe-oF, can help us understand how we can challenge this assumption in an informed and meaningful way.

Why did SAN fit?

Whether for content capture, post-production editing, visual effects, or coloring and finishing, Fiber Channel and Storage Area Networks met the needs of the modern media production environment for one key reason; they provided a shared storage environment that guaranteed video content could be played back without disruption.

But this capability comes at a cost, both in monetary terms, and the fact that a SAN is an 'island' of data that limits collaboration, productivity and efficiency.

“provide a shared storage environment that guaranteed video content could be played back without disruption”

How does SAN do it?

At its simplest, a Storage Area Network or SAN is a storage solution that provides deterministic connectivity between Hosts (servers and workstations) and Storage (disk or RAID devices), via a fabric. Each workstation, or host, within the fabric could address shared storage as if it were a local drive directly connected to the Operating system.

Because of its low latency and high bandwidth, it provided some degree of reassurance that the shared storage would not be a limiting factor in achieving video playback. However, the cost of SAN is often prohibitively high, requiring budgets and engineers to be dedicated to specific parts of the workflow. At a technical level, SAN requires additional ethernet networks to manage file locking, space allocation and data access and larger SAN fabrics can lead to unpredictable performance.

The overall infrastructure architecture and business direction may need to make compromises imposed by having to share data between services via a SAN, which may not be realised until after large investments have already been made. The overall cost and complexity is very high.

Where commodity IT fell down

Thanks to advances in the performance of commodity IT hardware and Ethernet networks, there was the potential to meet the performance requirements for Media & Entertainment with reduced complexity and vastly improved value for money per port or per client.

However, Scale-Out NAS Appliances and Software-Defined Data Platforms are not capable of servicing true 4K frame-based playback requirements. Even with the most careful tuning and the use of local SSD or NVMe cache on their nodes, they just cannot guarantee performance to real-time applications.

This is because the low level data-flows between NAS and SAN are dramatically different. The host Operating System must address a network storage target as opposed to a direct attached volume. This affects the application performance, and adds additional CPU overhead on both host and server side.

pixstor

The pixstor solution from pixitmedia is a holistic storage solution specifically designed for Media & Entertainment.

Software-defined, utilising commodity IT and Ethernet technology and taking a top down, application-first approach, pixstor achieves the flexibility of a shared NAS environment with the option of utilising technologies such as NVMeOF to provide a 'SAN-like' experience to specific hosts and applications. Utilising technologies such as NVMeOF to provide performance and latency that exceeds that of a SAN and gives that SAN-Like' experience".

At all times the pixstor solution respects the golden rule in Media & Entertainment; guarantee bandwidth to real-time video playback applications.

On-Set Dailies

Handling media content on-set is a punishing task for shared storage both in terms of real time playback requirements but also the number of concurrent read and write processes taking place simultaneously.

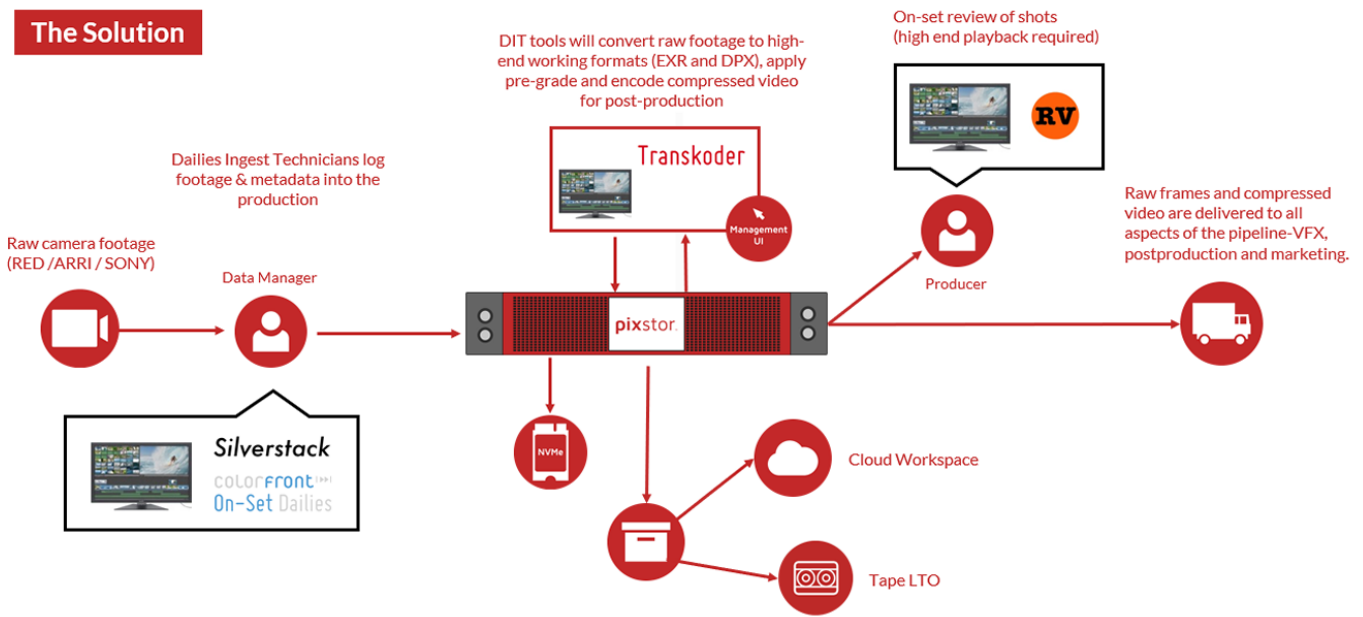
Camera footage is being constantly written to the file system, transcode processes are taking place to transform raw footage into frame-based and

“the overall cost overall cost and complexity of SAN is often prohibitively high”

“Scale-Out NAS Appliances and Software-Defined Data Platforms are not capable of servicing true 4K frame-based playback requirements”

“respect the golden rule in Media & Entertainment; guarantee bandwidth to real-time video playback applications.”

The Solution



compressed video files used in post-production and vfx pipelines, plus data is being written out to a local archive and delivered out to the many in-house and third party vendors downstream.

There is also a need for production staff to review content at its highest quality, in real time, at any point using software such as Autodesk Shotgrid or Colorfront Express Dailies.

Colorfront Software

Colorfront Transkoder and Express Dailies are applications that have optimised their application stack to be able to leverage both SAN and Ethernet shared storage. With multiple GPU systems capable of encoding raw video at faster than real-time speeds, and the ability to play back frame-based formats such as 16bit EXR, the application can comfortably exceed the limits of what 8Gb, 16Gb and even 32Gb Fiber Channel can deliver.

From benchmarks performed in the pixitmedia lab, we have proven that a pixstor solution with ATTO 100Gb NVMe-oF connectivity is capable of far exceeding SAN performance. And from a cost and complexity perspective, even the equivalent 25GbE pixstor and ATTO solution is less complicated and far more efficient than the equivalent SAN solution.


“we have proven that a pixstor solution with ATTO 100Gb NVMe-oF connectivity is capable of far exceeding SAN performance”

Benchmark Pipeline

The benchmark pipeline test is performed with Transkoder renders. Material is simply read and displayed as quickly as it can be, which provides a good sense of the ‘read’ performance of the underlying storage. These numbers should be considered the maximum that the end to end setup is capable of (including the workstation).

“less complicated and far more efficient than the equivalent SAN solution”

Transkoder Benchmark Pipeline




*Real-world achievable fiber channel bit rates and fps will be lower.

Source	Observed FPS	Observed Bandwidth (Gbps)
4K Scope 16Bit TIFF	177	56
4K Full 10Bit DPX	177	64
8192x4320 10Bit DPX	32	34.4

Real-World Results

The following test involves using Transkoder renders to convert a set of image sequences to video and back again, using different resolutions. The formats that have been chosen to test with are known to be fairly taxing from a CPU, Storage and GPU perspective, but without being overly 'brutal' on any one particular area. They are also the formats chosen by customers who have asked us to test Transkoder performance against pixstor in the past.

Transkoder Real World Tests



Source	Destination	Observed FPS with NVMe-oF (RDMA/Block)	Observed Bandwidth (Gbps)
4K Scope 16bit TIFF	4K Scope DCI (MXF/JPG2K)	60	19.2
4K Scope DCI (MXF/JPG2K)	4K Scope 16bit TIFF	55	17.6
2K Scope 16Bit TIFF	2K Scope DCI (MXF/JPG2K)	155	13.6
2K Scope DCI (MXF/JPG2K)	2K Scope 16Bit TIFF	165	14.4

**Tests were performed using a HP Z8 G4 workstation, 2 X Intel Xeon Gold 6136 CPU @ 3GHz (HT Enabled, 24 cores/48 logical) with 192GB Memory. The Compute GPU was 1x Nvidia GeForce GTX 1080Ti and the network interface was a ATTO FastFrame N312 Dual Port 100Gb Adapter.*

Summary

Application performance and ultimately end-user experience is affected by the whole solution stack; from the workstation, the access protocols in use, through to the network environment, shared filesystem and underlying disk configuration.

With this understanding, if we reframe the question of 'what is fit for purpose' then we can ensure that we have the right solution for any possible sector, workflow or pipeline within the Media & Entertainment industry, regardless of performance requirements.

The many moving parts within the modern production environment means we can no longer be lulled into the false sense of security that SAN technology used to provide.

It is with the pixstor solution, combined with ATTO technology, and our top down approach that we achieve the most effective way to ensure that the solution will be fit for purpose and we can focus on the job in hand; creativity and excellence.

“the pixstor and ATTO 360 solution is 100% cloud compatible”

“focus on the job in hand; creativity and excellence”

About pixitmedia.

pixitmedia delivers seamless collaboration to enable the power of ideas.

Our purpose-built, software-defined storage and data solutions simplify the flow of data to connect an increasingly complex world.

Our aim is to deliver beyond expectations throughout all areas of our operation. We devise solutions that give customers both choice and freedom, our restless innovation constantly pushes boundaries and the unrivalled care and knowledge of our team ensure optimum performance and value. Customer success is at the heart of our business.

We have a dedicated in-house lab facility to guarantee the effectiveness of our solutions.

pixitmedia is privately-owned and headquartered in the UK, with offices in the USA and Germany.

For more information on pixitmedia solutions:

UK Sales: +44 (0) 345 052 3721

e: info@pixitmedia.com

US Sales: +1 (424) 537-4948

w: pixitmedia.com