NVMe In the Studio; For 8K and Beyond

Pixit Media and Excelero
8K+ Workflow Solution Brief

Highlights

Blazingly Fast
Bypass storage server bottlenecks to deliver the highest possible performance, accelerating uncompressed 8K and beyond workflows

Maximum Throughput to a Single Client
Over 20 GB/s single stream performance to a workstation or server over Ethernet

Performance Scales Linearly
Aggregate performance scales linearly by simply adding more NVMesh target nodes and NVMe drives

Software Defined Future-Proof Architecture
Studios can keep up with ever-expanding resolution, frame-rate and color depths though incremental improvements in commodity server, network and NVMe storage technology over time "for free", without having to disruptively tear out proprietary hardware and start again

Super Simple Hardware
PixStor and NVMesh use off the shelf commodity hardware with no bespoke components

Near Zero Overhead
There is almost zero overheard on either the client or NVMesh target server delivering the performance to the client, leaving CPU and memory almost 100% available for applications

Low Power and Space Needs
Performance rates as high as 40 GB/s per U of rack space, important in M+E environments where space and power are often expensive and in short supply

Easy Setup and Management
Setting up NVMesh is simple with automated deployment as part of the PixStor stack and an intuitive GUI enhances the management experience

The High-End Editing, Finishing, and Visual Effects Data Challenge
Moore's law has overtaken modern film and video production in multiple dimensions. Resolution, frame-rate and color depth are all on the march simultaneously. The arms race for clearer, crisper, more lifelike content is here and whatever image quality producers demand today is guaranteed to be obsolete tomorrow – driving the demand to create content now that will be future proof for years to come. The revolution in high-resolution uncompressed workflows has overwhelmed the throughput capabilities of legacy Fibre Channel and Scale-Out NAS storage, making it increasingly unaffordable to use traditional RAID-based storage hardware to deliver performance to workstations and servers. Fortunately, a simultaneous revolution is occurring in storage – delivering new levels of breakthrough performance.

NVMe in the Studio with Pixit Media PixStor and Excelero NVMesh®
That revolution is NVMe, the astonishingly fast interface protocol that bypasses legacy storage bottlenecks, getting data directly to servers and high-end workstations with no speed bumps. But left on their own, local workstation NVMe drives are under-utilized, non-redundant and individually managed. Pixit Media PixStor with Excelero NVMesh takes NVMe drives in individual servers, or dedicated storage servers, and virtualizes them into aggregated volumes that can be rapidly accessed as a shared resource by multiple clients over simple Ethernet. Excelero NVMesh with PixStor uniquely delivers centralized management, redundancy and high utilization rates, while maintaining local NVMe performance.

Editing, finishing and VFX applications can now enjoy the low latency, high throughput and IOPS of a local NVMe device while getting the benefits of centralized, redundant storage. The combined Pixit Media PixStor and Excelero NVMesh solution provides a single namespace for applications that demand extreme performance. As shown below, a near-line tier can be added using traditional storage arrays with PixStor's data placement policy engine to automatically tier data to lower cost storage for high capacity requirements – with workflows extending to the cloud.

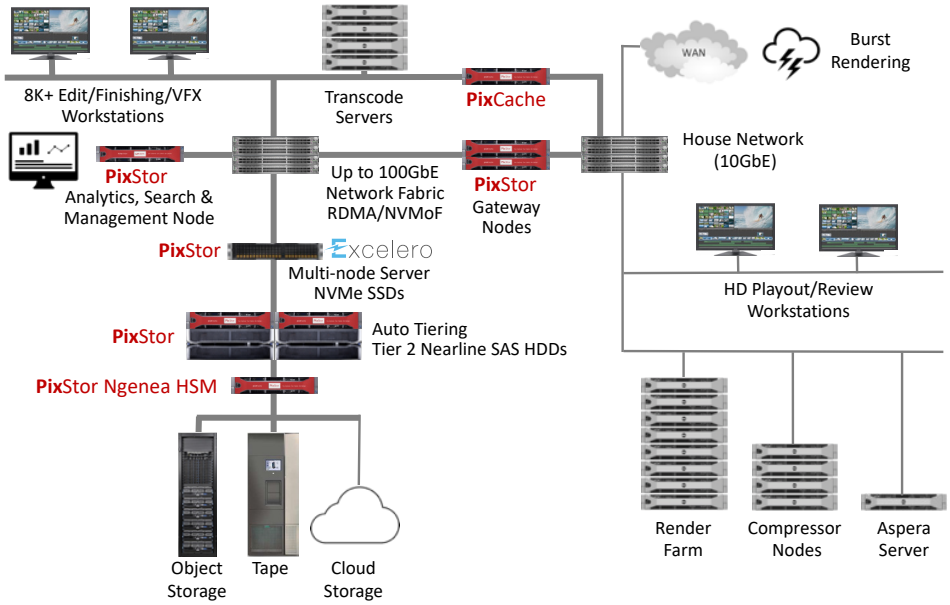

Pixit Media Excelero Single Namespace 8K+ Cloud Integrated Architecture

pixitmedia

Excelero

Advanced PixStor Software

- Centralized control of all file system components from a single interface
- Pythonic and REST-compliant toolset can be used to automate Spectrum Scale operation and create customized web interfaces
- Vastly reduced complexity for otherwise difficult to manage file system components
- Monitoring, alerting, search and analytics
- High performance tiering to object storage with full path retention so object is directly accessible by cloud applications, enhancing workflow
- Seamless cloudbursting
- Cost effective, simplified Disaster Recovery via remote PixStor asynchronous replication

Unified Services and Support

- "Disk to desktop" holistic solution understanding
- Direct access to knowledgeable expertise – no "briar patch" support
- Hybrid managed service: regular maintenance tasks such as firmware and upgrades are handled by PixStor support, rather than the customer's technical staff
- Alerts generated by the system are proactively investigated and escalated by PixStor support
- Single point of contact multi-vendor incident management ensures SLA's and top-quality support experience is delivered
- Automated weekly health checks

Comprehensive Solution Design

- Extensive, in-depth analysis of existing environment focusing on workflow requirements, storage, network, and applications
- Identify issues and bottlenecks
- Design solution architecture including hardware, software configuration and procurement guidance
- Integrated system build and implementation

PixStor with Excelero Supports More 8K+ Workflows than any Other Solution

Pixit Media PixStor is a high performance, highly scalable, enterprise-class storage solution specifically designed for Media & Entertainment workflows. PixStor combines flash, disk, tape and cloud storage with affordable, high performance Ethernet into a unified system that's higher performing, limitless in scale and lower cost than traditional legacy solutions. Data moves seamlessly through many tiers of storage – from fast NVMe SSD to cost-effective, high capacity object storage, all the way out to the cloud – depending on how frequently it needs to be accessed. This allows media organizations to accelerate high resolution workflows and store valuable assets more safely and economically.

Removing Legacy Bottlenecks Delivers Extreme Performance

PixStor Excelero NVMesh volumes are served out over 100 Gb Ethernet, through standard 100 GbE Mellanox Spectrum Ethernet switches. With Excelero NVMesh, there are no bottlenecks from RAID controllers or proprietary hardware. With all those legacy layers removed, all that remains is the raw speed of the NVMe SSD, the speed of the PCI bus on the server and the speed of the network connection. As these industry standard technologies improve, both PixStor and Excelero can exploit immediate performance boosts with small, incremental upgrades. Incremental improvements to commodity server, network and storage technology means incremental improvements to NVMesh performance 'for free', without having to tear up everything and begin again.

Scale to 1000's of NVMe Drives with Virtually No CPU Overhead

High end workstations connected to the NVMe network fabric see and address volumes directly, without having to go through the storage servers. If a NVMesh node or drive fails, redundancy is handled transparently without the need for administrative intervention. What you see on the clients is not the actual NVMe drives, it is the virtualized volumes. Tens of thousands of NVMe drives can scale capacity and performance linearly within a single namespace, making management and usage easy. Unique to combined PixStor and Excelero solution, there is almost zero overheard on either the client or NVMesh target server delivering the performance to the client, leaving CPU and memory almost 100% available for applications.

Superior Performance on Simple Industry Standard Hardware

Each NVMesh target can deliver 20 GB/s of read performance. An NVMesh target node fits within a simple, commodity server configuration that contains a CPU, a relatively small amount of memory, NVMe drives and Ethernet connectivity. Off-the-shelf, Industry standard servers can fit four of these nodes in a 2U chassis, delivering over 80 GB/s of read performance to clients. This reduces power and space requirements by an order of magnitude, compared to conventional disk-based storage. This is an important benefit in many modern Studio environments where space and power are both very expensive and in very short supply.

Blazingly Fast Performance to a Client

Pixit Media has measured over 20 GB/s to a single server or workstation from an NVMesh system using only 2 X 100 Gb Ethernet connection. Comparably, throughput requirements under 10 GB/s can be delivered with only a single cable and card. With traditional Fibre-Channel based solutions, up to 8 X 16 Gb FC connections would be required to meet this performance, with a much lower value per port compared with 100 Gb Ethernet.

High Performance Parallel File System with Limitless Scale

PixStor is powered by IBM Spectrum Scale enhanced with added ease of use, collaboration, tiering, cloud integration and relentless support – on open standard hardware – greatly reducing costs

Transparent Data Tiering to Object, Tape, and the Cloud

Reserves high speed storage for work in progress, moves everything else to low cost archive, while maintaining a single view of all data

Effortless Burst Rendering

Painlessly cloudburst onto unlimited nodes, minimizing data transfer and maximizing data locality

Simplified Management and Workflow Automation

Easy Python API programming interface removes complexity and provides 'DevOps' style management capabilities for systems administrators and pipeline developers – automating management and workflows

Complete Data Awareness and Lightning Fast Search

Intuitive user defined metadata harvesting, tagging and easy search for data – no more tree walks or lost data

Advanced Analytics

Real time performance, history and trending with statistics reporting about data and its contents through a simple and intuitive graphical interface

Multi-protocol

SMBv2/3, NFS, FTP, HTTP, S3, HDFS and POSIX parallel file system client for extreme performance

Linearly Scalable Aggregate Performance

Aggregate performance scales linearly by simply adding more NVMesh target nodes and NVMe drives. The table below shows what can be achieved in terms of stream counts from a typical 3 NVMesh target starter configuration (60 GB/s), which can fit in as little as 2U of rack space.

Format/FPS	Aggregate Stream Count
10 Bit DPX HD, 24 FPS	270 streams
10 Bit DPX 2K, 24 FPS	180 streams
10 Bit DPX, 4096 x 2160, 24 FPS	75 streams
16 Bit EXR, 4096 x 2160, 24 FPS	48 streams
10 Bit DPX, Full 4K, 24 FPS	48 streams
16 Bit EXR, Full 4K, 24 FPS	34 streams
10 Bit DPX, 8192 x 4320, 24 FPS	18 streams
10 Bit DPX, 8192 x 6224, 24 FPS	12 streams

3 NVMesh Target Server Configuration Stream Count

Easy Set Up and Management

Setting up NVMesh is simple and is automated as part of the overall PixStor deployment. The GUI is very intuitive and Excelero has a REST API that can be used for DevOps automation.

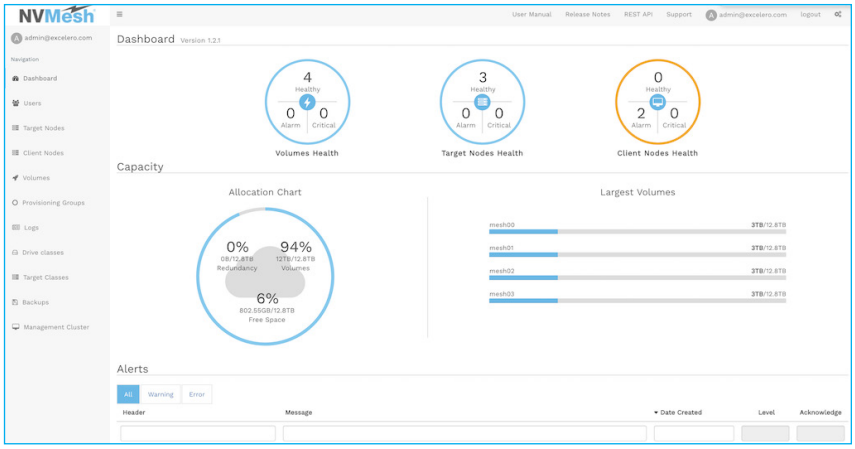

Excelero NVMesh Management Dashboard

A Completely Software Defined Future-Proof Architecture

PixStor with Excelero is deployed on software-defined, open standard hardware, with no vendor lock-in, and includes most data protection services at no extra charge – providing media organizations economic advantage and more purchasing power, now and in the future. Being completely software defined, Studios investing in the Pixit Media Excelero solution can keep up with ever-expanding resolution, frame-rate and color depths though incremental improvements in commodity server, network and NVMe storage technology over time, without having to disruptively tear out proprietary hardware and start over.

pixitmedia

Excelero

Compute Resources Go to Applications, Not Storage

Other NVMe solutions force the computation overhead of the volume placement on the client itself, providing lots of performance, but little resource left for applications to do anything with that performance

Uses Mellanox Spectrum 25/50/100G Ethernet Switches

Mellanox switches form a network fabric that unleash the maximal power of the NVMe storage pool. Powered by Mellanox's own ASICs, they deliver non-blocking, cut-through switching at line speeds with no packet loss. The dynamically shared switch buffer provides the best microburst adoption enabling PixStor with Excelero to deliver guaranteed throughput and latency.

Data Placement Policy Engine

Unique to PixStor with Excelero is the ability to intelligently tier data to lower cost storage automatically based on policy

Readily Support Future Workflows

Be ready for the future by harnessing the power of a hardware agnostic strategy – driving down costs and enabling enduring flexibility

No More Barriers

From a reliability and performance perspective, the shackles are off, PixStor with NVMesh has removed all barriers to delivering performance at scale

Combining all Workflows in a Single Namespace Tightly Integrated with the Cloud

PixStor with Excelero provides the performance needed so combined VFX, finishing and editing workloads can all run at maximum speed without effecting editors' or artists' user experience. What's more, in addition to tiering data to the cloud with PixStor Ngenea HSM, PixStor PixCache facilitates efficient Burst Rendering onto unlimited compute nodes to take on extra projects or accelerate schedules. PixCache minimizes data transfer and maximizes data locality by caching on premises data close to cloud computing resources. Results are automatically transferred back to on-premises storage with no manual intervention.

An Industrial Strength, Robust, Easy to Manage Storage Platform

PixStor is an easy to manage, scalable, high performance storage system with centralized control of all file system components from a single interface employing graphical interfaces for common tasks, and Python and REST APIs for file system automation and workflow integration. PixStor monitoring, alerting and performance analytics are tailored to the exact needs of the administrative workflow and empower the organization with unprecedented insight into data usage and trends.

Under the Hood: What makes PixStor So Capable

PixStor is powered by IBM Spectrum Scale, a proven high-performance parallel file system trusted by thousands of organizations worldwide. IBM Spectrum Scale can easily manage petabytes of data and billions of files, all under a single global namespace. IBM Spectrum Scale stripes data in parallel to multiple disks across multiple servers delivering extreme, scalable performance – and supports the newest low-latency, high-bandwidth NVMe flash technology. IBM Spectrum Scale scatters data across spinning disk, removing the impact of fragmentation on performance, so there is no degradation in performance as the system fills – as is common in many file systems. The performance of the system enables multiple workloads to run full out with no negative effect on finishing editors' or artists' user experience.

All-inclusive, Expertly Delivered "Disk to Desktop" Unified Services and Support

We uniquely support our customers' complete environment, providing direct access to knowledgeable expertise and hybrid managed services done by Pixit Media engineers.

Pixit Media has Become the De Facto Standard for Post and Broadcast in the UK

Pixit Media solutions are deployed by leading content creation and distribution organizations around the world. We have been counted on to deliver numerous commercially successful Oscar, BAFTA and Emmy winning TV and Film VFX and CGI, Commercials and and Music Videos.

Example Pixit Media Customers in Post and Broadcast

pixitmedia

For more information on all Pixit Media Solutions:
T: +44 (0)845 052 3721 | E: info@pixitmedia.com

Excelero