

Unleash Your Storage Power with NVMe-oF on All-NVMe Servers

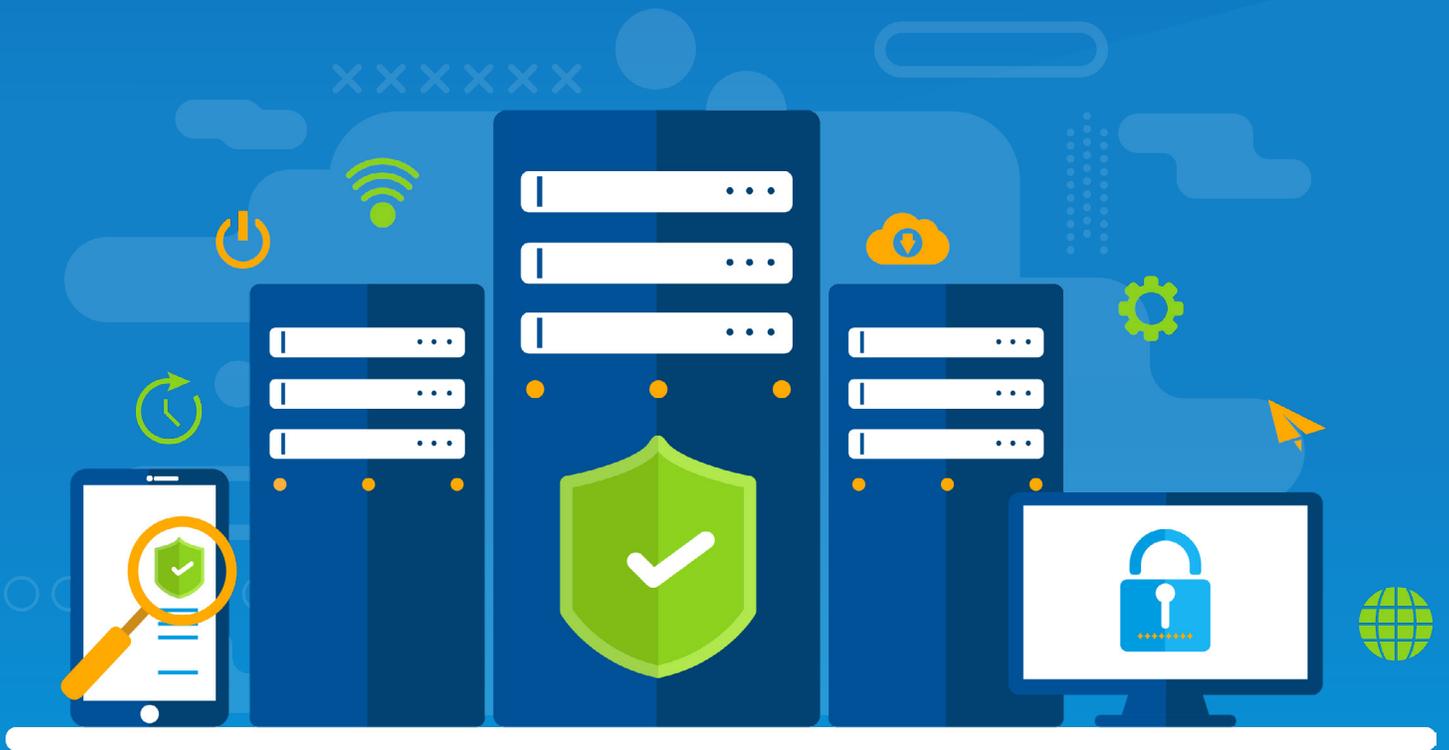




Table of Contents

Introduction	3
Achieving High Storage Performance Levels with NVMe-oF	4
NVMe-oF Deployment Success: Excelero NVMesh on QCT All-NVMe Storage Servers	5
NVMe Use Cases Featuring QxSTOR Excelero NVMesh	9
Conclusion	11

Introduction

Today, new workloads like AI/ML and real-time data analytics redefine the IT infrastructure. The need for higher throughput and lower latency is now all the rage. However, legacy infrastructure may fail to deliver desired performance.

In this whitepaper, we look at ways to achieve greater storage efficiency and performance by deploying All-NVMe servers for low latency and high storage performance in evolving workloads. We will explore what NVMe over fabrics (NVMe-oF) benefits offers for performance-sensitive use cases and how QCT and Excelero's integrated solution rivals high storage performance.



Achieving High Storage Performance Levels with NVMe-oF

Increasingly, IT leaders recognize the power of flash storage to provide high throughput and low latencies. To meet these needs, the NVMe protocol was extended across various network fabrics (e.g., Fibre Channel, Ethernet and Infiniband) to reduce data handling overhead and latencies as well as eliminating bottlenecks. These standards ensure fast connectivity to handle massive data and the demands of newer workloads, such as AI and machine learning. IT administrators can deploy NVMe-oF storage to rival direct attached storage (DAS), making it possible to share flash devices among servers as needed. For instance, organizations regularly encounter data bottlenecks when they are using GPU servers to run data-intensive AI and ML workloads. By implementing NVMe-oF over a high-speed network can enable these servers to access unlimited pools of shared NVMe storage and gain the same performance characteristics as local attached NVMe SSDs.

However, it could be tedious for IT administrators to set up NVMe-oF storage and monitor its status. QCT and Exelero jointly built an easy deployment and monitor solution while scaling with predictable performance.





NVMe-oF Deployment Success: Excelero NVMesh on QCT All-NVMe Storage Servers

Despite the advantages of implementing an NVMe-oF approach to storage, successful adoption requires committing substantial IT resources. For example, it's necessary to maintain separate server and storage resources (disaggregation) to achieve targeted performance goals as well as reduce storage controller latency overhead. IT teams may face service interruptions when they try to configure existing NAS or SAN systems during NVMe-oF deployments. Administrators need to spend valuable time provisioning and monitoring NVMe resources using the traditional command line interface in addition to monitoring SSD health, updating firmware, ensuring software compatibility and performing other tasks.

Overcoming these limitations, Excelero NVMesh simplifies NVMe-oF deployment across traditional network fabrics to support local or distributed file systems. Built to provide customers maximum flexibility in designing their storage infrastructures, NVMesh supports not only traditional Fibre Channel and Ethernet but also cutting-edge fabric networks like InfiniBand and RoCE.

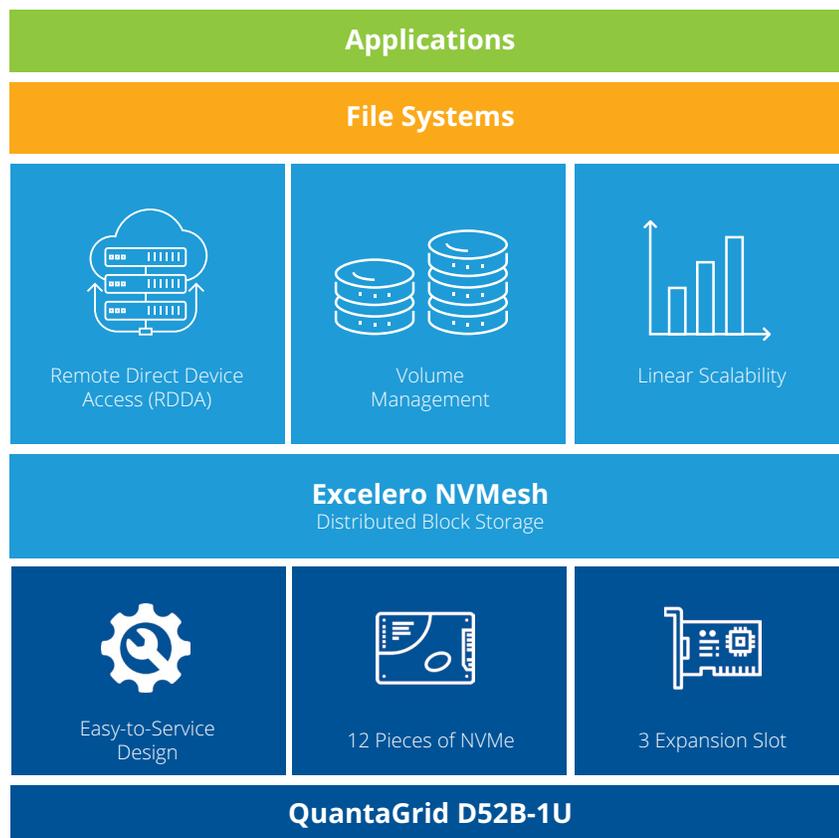


Figure 1: High level QxStor Excelero NVMesh stack



NVMe is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures for efficient pooling of NVMe-based storage using NVMe-capable servers. As software-designed block storage, NVMe delivers the necessary low latencies and fast response times that today's applications require while providing the scalability and management of distributed drives. The Excelero solution scales granularly as a single pool of high-performance block storage, enabling IT administrators to easily provision volumes and ensure that their applications can access block volumes directly via file systems.

Excelero NVMe has been performance benchmarked to demonstrate true linear scalability, which refers to the ability to maintain a consistent, proportional throughput rate as resources are added to the cluster. As deployment requirements change, IT teams can add or remove cluster nodes as needed. In general, true linear scaling can be difficult to achieve because deploying new resources incurs additional overhead, but NVMe deployed on NVMe-capable servers can scale up to 128-node clusters with minimum degradation.

Administrators consistently face challenges when they try to deliver new levels of server and storage performance to meet application demands, particularly in the areas of ML, AI and Big Data analytics. Pooling local high-performance storage, such as SSDs, has required IT to use either valuable CPU cycles to manage storage across the network or compromise on overall performance.

Excelero's patented remote direct drive access (RDDA) functionality incorporates the RDMA protocol to eliminate these constraints and allow data access with no remote node CPU involvement. Bypassing the CPU to read and write data directly to NVMe drives offers two key benefits:

1. It removes the “noisy neighbor” problem where the performance of one application can impact another.
2. RDDA eliminates any effect on target CPU performance.

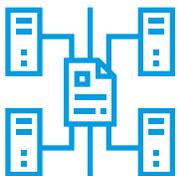
From Quanta Cloud Technology (QCT) — a subsidiary of Quanta Computer, Inc., a Fortune Global 500 technology engineering and manufacturing company — comes the All-NVMe Storage Server (D52B-1U) with Excelero NVMe, an integrated solution is designed for a variety of performance sensitive workloads like database or HPC that regularly consume expensive compute resources and memory, while requiring fast I/O performance.



Powered by the latest Intel® Xeon® Scalable Processors, this world-class high density 1U server supports up to 12 2.5" Intel® SSD DC P4510. With Intel's 64-layer, 3D NAND technology, Intel® SSD DC P4510 offers increased density to support broader workloads and improve data service levels. It delivers optimized performance and scalability, increasing server agility and utilization in cloud storage applications to minimize TCO. By utilizing Intel® Xeon® Scalable processors and Intel® SSD DC P4510 for high-performance storage and using the QCT All-NVMe Storage Server with Excelerio NVMeMesh, customers can achieve full design flexibility and gain unprecedented performance speeds for remote storage. QCT offers services for different types of scale-out storage clusters in a range of configurations to enable diverse workload optimizations for IOPS, throughput and cost/capacity.

Administrators can use the Excelerio NVMeMesh's web-based GUI and feature-rich performance analytics engine to pinpoint anomalies quickly and facilitate maintenance. They can also access cluster-wide and per-object utilization statistics to monitor and analyze performance at scale. As a result, IT teams can leverage widely deployed NVMe-based SSDs, treating them as unified, redundant storage pools and avoiding the latencies that occur when accessing distributed storage across a network.





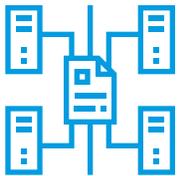
NVMe Use Cases Featuring QxSTOR Excelero NVMesh

Use Case #1: Accessing Parallel File Systems with NVMesh

As HPC deployments push the limits of available resources, NVMe-oF storage provides an essential foundation for executing a greater number of concurrent, parallel I/O operations while off-loading CPUs to address more actual compute tasks. For example, parallel file systems store data across multiple networked servers and facilitate high-performance access through simultaneous I/O operations between clients and storage nodes.

However, it's complicated and time-consuming for IT teams to configure NVMe-oF cluster by themselves. In contrast, administrators deploying Excelero NVMesh can quickly scale NVMe storage capacity and bandwidth to process enormous amounts of data with features that include high availability, mirroring, replication, while utilizing centralized management, monitoring and administration. Providing fast, elastic, pooled NVMe storage beneath a parallel file system improves metadata handling and ensures a higher read performance and shorter latencies resulting in higher storage utilization.

Mainstream HPC applications and AI/ML typically have extremely low latency and high I/O throughput requirements. These workloads and applications comprise both large blocks of data and metadata (consisting of numerous small files). For example, a big data analytics workload would rely on one massive dataset while an ML application might use millions of small files. Excelero NVMesh offers MeshProtect to enhance metadata server performance compared to traditional software RAID by parallel file system itself. Its parity-based data redundancy scheme runs on the clients in a decentralized fashion to help achieve near-zero CPU utilization and free up CPU cycles for parity calculations.



Use Case #2: Capitalizing on GPU Servers for AI and ML

Graphics processing units (GPUs) are specialized processor cores that speed computations and can be used effectively in parallel for calculations. These GPU cores were originally designed to process images and visual data, but they're now widely deployed for more advanced AI and Deep Learning applications.

Multiple AI deployments and ML training datasets demand significant consumption of infrastructure resources: ML datasets size can far exceed a single GPU's local RAM capacity, creating an I/O chokepoint and leading to rampant storage bottlenecks. As a result, AI applications and ML workloads under-perform and encounter unacceptable latencies due to shortfalls in GPU responsiveness and storage limitations.

Local NVMe flash represents the only resource that's fast enough to keep up with specialized GPUs and eliminate GPU storage bottlenecks. However, AI/ML might require scalable and shared NVMe storage when data sets grows larger. Excelero's NVMeMesh bridges the divide between performance and practicality, ensuring that GPU servers can access scalable, high performance NVMe storage pools as if they were local. Organizations and IT leaders gain the benefits of higher ROI, easier workflow management and faster time to results.



Conclusion

Innovations in IT hardware, networking and software are revolutionizing the delivery of storage systems to the enterprise. Comparable strides in processor design are enabling modern applications to reach new performance possibilities, particularly through increasing reliance on All-NVMe servers combined with NVMe-oF deployments.

To be sure, the storage I/O capabilities of flash optimized with PCIe enable modern multi-core processors to complete computations faster, relaying more data with less utilization of processor CPU time. These latest storage standards have been at the forefront for ensuring fast connectivity to handle the massive data increases and demands of AI, data analytics, ML and comparable technologies.

One of the primary benefits of NVMe is to enable true converged infrastructure by logically disaggregating storage from compute. To achieve these goals, QCT has partnered with Exceero to meet today's storage requirements and maximize data center ROI by ensuring unlimited data feeds. Through Exceero NVMe and All-NVMe QCT Storage Server deployments, customers can combine the power of the fastest parallel file system with the fastest block storage. This end-to-end scale-out, high-performance storage solution delivers impressive performance benefits to mission-critical workloads across industries.

QCT offers a full spectrum of data center products and services from engineering, integration and validation to global supply chain support. QCT not only provides data center products but also offers cloud solution to cloud service providers and large enterprises, and through channel partners to enterprises and SMBs. Through collaboration with various cloud solution ISVs, QCT equipped themselves with solid knowledge for software-defined data center solutions, with a proven track record in delivering software-defined storage and HCI solutions.

To learn more about the benefits of NVMe and All-NVMe QCT Servers, go to <https://go.qct.io/solutions/software-defined-storage/qxstor-exceero-nvme-solution/>