


*Excelero delivers low-latency distributed block storage for high-performance workloads. Founded in 2014 by a team of storage veterans and inspired by the Tech Giants' shared-nothing architectures for web-scale applications, the company has designed storage software that facilitates Elastic NVMe for AI training, data analytics and containers. Excelero's NVMesh enables shared NVMe across any network and supports any local or distributed file system. Customers benefit from the performance of local flash with the convenience of centralized storage while avoiding proprietary hardware lock-in and reducing the overall storage TCO. NVMesh delivers 20x faster data processing for multi-server, multi-GPU compute nodes when working with massive datasets for complex financial analysis, climate modeling and other HPC workloads.*

Analytics capabilities are heavily influenced by the volume of data that can be analyzed and the speed at which this can be done. This makes storage infrastructures the backbone of next-generation analytics applications. It doesn't matter how powerful the analytics engine is, or how feature-rich the used applications are if the storage infrastructure does not meet performance and scalability requirements.

Just like GPUs are reshaping high-performance computing, NVMe flash is revolutionizing storage performance. The latest generation of NVMe flash devices can deliver over one million read IO/s, over 6GB/s of bandwidth and under 20µs of latency. But the problem is that this performance can only be consumed locally and GPUs can process data much faster than what local NVMe devices can feed them. To solve this data bottleneck, GPUs need access to unlimited pools of shared NVMe with the same performance characteristics as local NVMe: Elastic NVMe.



## NVMesh Elastic NVMe

Excelero delivers low-latency distributed block storage for high-performance workloads. NVMesh is storage software that enables Elastic NVMe for AI training, data analytics and containers. The solution features an intelligent management layer that abstracts underlying hardware with CPU offload, creates logical volumes with redundancy, and provides centralized management and monitoring. Customers benefit from the performance of local flash with the convenience of centralized storage while avoiding proprietary hardware lock-in and reducing the overall storage TCO. NVMesh delivers 20x faster data processing for multi-server, multi-GPU compute nodes when working with massive datasets for complex financial analysis, climate modeling and other HPC workloads.

NVMesh is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures, giving customers full freedom in their architectural design. Leading customers include The Science and Technology Facilities Council (UK), Instadeep (UK), Bezeq (IL), Deutsche Bank (Ger), Barclays (UK), NYC DOH (US), CCC (US) and Scinet (CA)

### PERFORMANCE AT ANY SCALE

Excelero's Mission is to provide the highest storage performance at any scale, without trade-offs on operational and financial efficiency.

It's our mission to enable our partners, customers and users to meet all their business goals:

**For our partners, we designed the most agile storage software to build end-to-end storage solutions for their customers' workloads:**



**Software-Defined Software:**  
server-based architecture,  
use any hardware



**Distributed Block Storage:**  
use any file system,  
scale-up or scale-out



**Elastic NVMe:**  
Leverage the fastest storage  
devices across the data center

For customers, we designed software-defined block storage that meets the performance and scalability requirements of their applications and the agility and efficiency requirements of their business. We enable customers to massively improve their ROI across the entire datacenter: use standard HW components, maximize component utilization (NVMe, GPU), minimize overhead (storage, networking, compute), reduce software licenses cost.

For end users, we designed a storage layer that eradicates data bottlenecks: access data at any speed in any location.

Excelero's markets are not defined by the business customers are in but the workloads they run. For those organizations that run high-performance workloads, we provide the highest-performance storage at the best price/performance at any scale.

NVMesh is software-defined storage, designed for partners to build end-to-end storage solutions for their customers' workloads

### WORKLOADS - USE CASES - APPLICATIONS:

- GPU Storage for AI – Machine Learning & Deep Learning Training
- Data Warehouse Storage (Analytical Workloads)
- Container Storage.

