

SOLUTION BRIEF

Intel® Storage Builders

Excelero NVMesh*

Intel® SSD DC P4600 Series

Distributed NVMe* Server SAN



Deploy NVMe*-Based Storage at Data Center Scale

Leverage Excelero NVMesh and Intel® SSD Data Center Family for PCIe* devices to maximize your NVMe ROI for web-scale applications.

“Excelero enables customers to leverage the full potential of NVMe at scale without giving up local performance characteristics, which lowers the storage cost dramatically.”

Lior Gal, CEO Excelero

Executive Summary

Current generations of servers are flexible systems, designed to enable Software Defined Infrastructure (SDI). They incorporate strong processing power, high-speed networking, local high-performance flash storage, and the latest-generation of memory. Advanced clustering methods enable pooling of CPU resources across multiple systems to tackle large compute jobs. However, pooling local high-performance storage media, such as Intel data center-class NVMe*-based SSDs, across the network has proven to be a challenge. IT must compromise on storage performance, use valuable CPU cycles for managing storage across the network instead of for applications, or both.

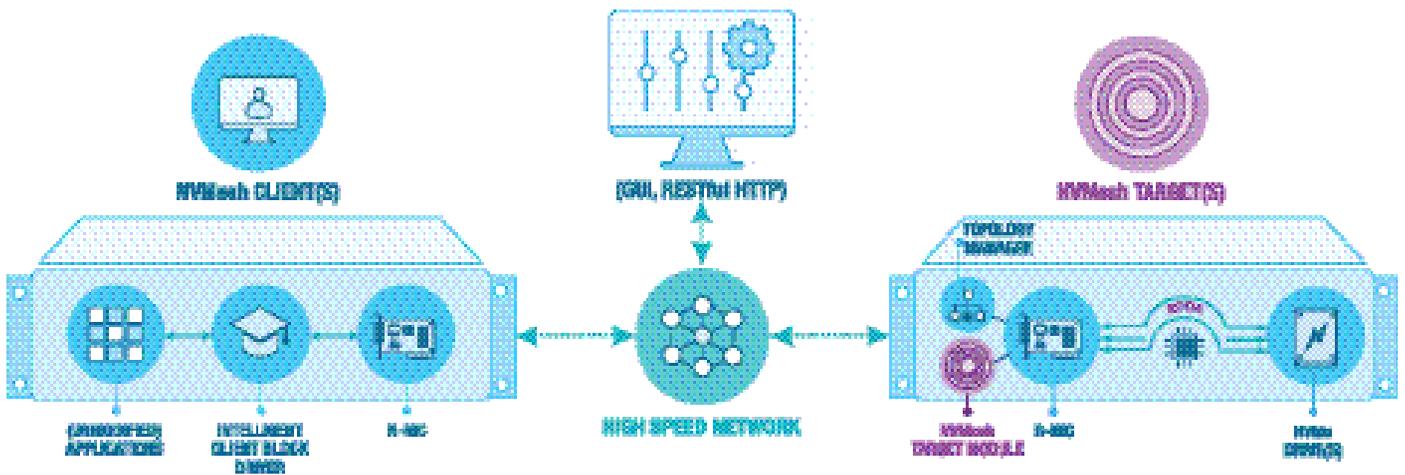
Excelero NVMesh* enables efficient pooling of NVMe-based storage, utilizing off-the-shelf components without impacting target CPU performance. By not requiring expensive proprietary hardware, it enables high-performance storage pooling and saves processor resources for their main purpose, the applications.

Excelero NVMesh – Enabling Distributed Storage with Local NVMe-Based SSD Performance

Excelero's NVMesh is a high performance, low-latency software-defined block storage solution designed to meet the performance and scalability requirements of next-generation applications. NVMesh bridges the speed and low latency of local storage with the scalability and management of distributed drives. It enables customers to deploy NVMe*-base storage devices, such as Intel® SSD Data Center Family for PCIe* devices, at massive scale—without adding typical network latency or modifying applications to manage data locality, all while using industry-standard hardware. NVMesh has been deployed for Industrial IoT, machine learning, and user behavior analytics applications.

Excelero NVMesh delivers in-server, direct-attached NVMe-based storage performance with data protection from a distributed storage infrastructure utilizing standard off-the-shelf components. It allows IT architects to leverage NVMe-based SSDs, potentially spread over many physical systems, and treat them as a unified, redundant storage pool without adding the inherent latency of distributed remote storage access over a network. Excelero NVMesh enables users to leverage the full performance of NVMe-based SSDs at massive scale—without modifying applications for data locality.





NVMe Architecture

The NVMe architecture features a distributed block layer that allows applications—without modification—to utilize pooled NVMe-based storage devices across a network at local speeds and latencies. Distributed resources are pooled with the ability to create arbitrary, dynamic block volumes that can be utilized by any host running the NVMe block client. These virtual volumes can be striped, mirrored, or both, while enjoying centralized management, monitoring and administration. Applications can enjoy the latency, throughput, and IOPs of a local NVMe-based storage device, while getting the benefits of centralized, redundant storage.

NVMe is deployed as a virtual, distributed, non-volatile array that supports both converged and disaggregated architectures, giving customers full freedom in their architectural design. An Excelerio solution scales granularly as a single pool of high-performance block storage, and it features easy provisioning of volumes. Applications can access the block volumes directly or through a file system, which makes integration very straightforward.

NVMe’s Remote Direct Drive Access

A key component of Excelerio’s NVMe is the Remote Direct Drive Access (RDDA) functionality, which bypasses the CPU. As a result, Excelerio NVMe avoids the noisy neighbors effect for the application. The shift of data services from centralized CPU-based to complete client-side distribution enables linear scalability, provides deterministic performance for applications, and enables customers to maximize the utilization of their flash drives.

Intel SSD Data Center Family for PCIe/NVMe Devices Deliver Data Center Reliability and Performance at Web Scale

The Intel SSD Data Center Family for PCIe brings extreme data throughput directly to Intel® Xeon® processors with up to six times faster data transfer speed than 6 Gbps SAS/SATA SSDs.¹ Intel led the industry in creation of a new Non-Volatile Memory Express* (NVMe) storage interface standard. NVMe overcomes SAS/SATA SSD performance limitations by optimizing hardware and software to take full advantage of NVMe SSD technology.

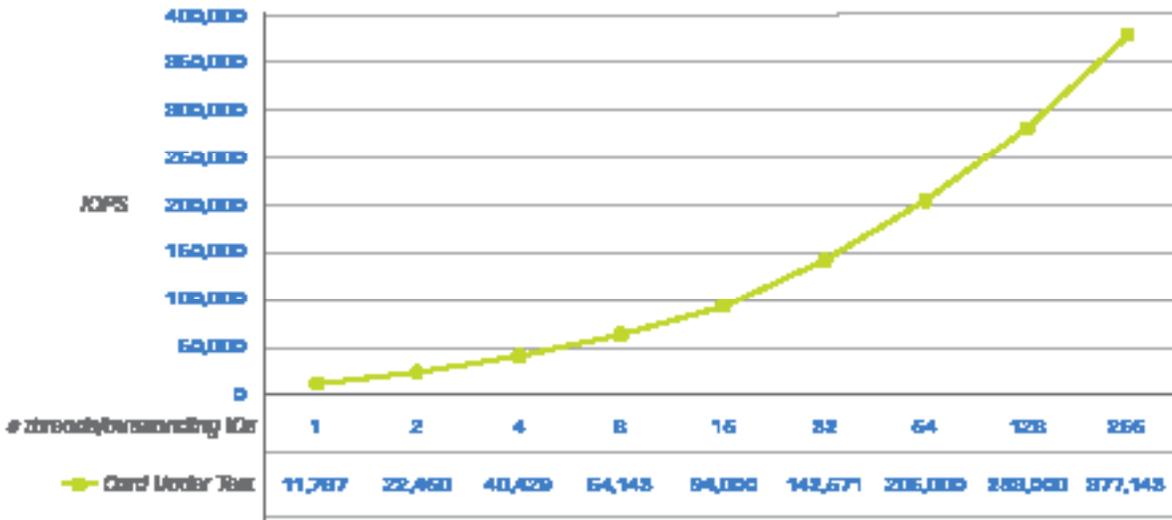
Intel SSD Data Center Family for PCIe devices are based on Intel-developed controller, firmware, and leading manufacturing process NAND flash memory. Rigorous qualification and compatibility testing ensures a highly reliable SSD. The Intel® SSD Data Center Tool provides a powerful set of management capabilities.

Recent testing by Excelerio revealed incredible performance using an NVMe client reading and writing a remote Intel® SSD DC P4600 drive using the Excelerio RDDA protocol. The transactions did not use CPU cycles on the target storage server.

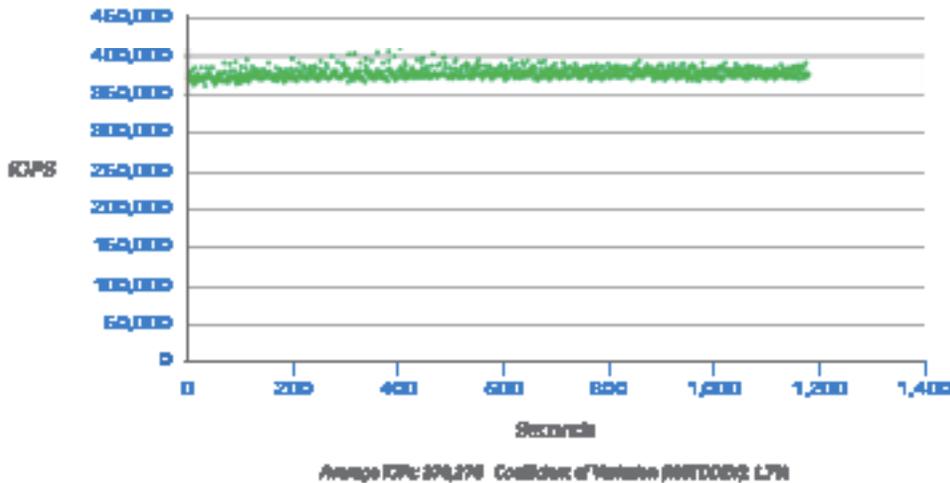
The following graphs show the mixed read/write performance and performance stability results from testing using the following configuration.

| Component | Specification |
|------------------|---|
| System Chassis | Supermicro* SYS-2028BT-HNR+ |
| Processor(s) | 2X Intel® Xeon® Processor E5-2620 v4 @ 2.10 GHz |
| Memory | 64 GB |
| Operating System | CenOS* Linux* 7.3.1611, kernel 3.10.0-514.16.1.el7.x86_64 |
| Network | Dual-port Mellanox* CX-5 @ 50 Gbps |
| Testing Software | Excelerio* NVMe ver. 1.2.0; FIO ver. 3.0.5-g168b |

Sustained 4KB Random Mixed 70%/30% Performance by # of Threads Using 100% Capacity



**Performance Stability
Sustained Multi-Threaded Random 4KB Mixed (w/30%/70%) using 100% Capacity**



storage requirements.

Use applications natively – NVMesh allows applications to utilize pooled NVMe-based storage—without modification—across a network at local speeds and latencies.

Simplicity – A unified pool of NVMe-based devices enables customers to maximize storage utilization and avoids data locality issues for the application.

Full convergence – NVMesh enables 100 percent converged infrastructure by full logical disaggregation of storage and compute.

Remote avoids the tions can

Enables NVMe-based SSD-based virtual flash is a highly flexible deployment model for server SAN. It is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures—and even mixed environments—giving customers full freedom in their architectural design.

Highly Scalable – NVMesh delivers a sharing technology that scales performance linearly at near 100% efficiency. The shift of data services from centralized CPU to complete client-side distribution provides deterministic performance for applications and enables customers to maximize the utilization of their flash drives.

IT Benefits

| Data Management & Protection | |
|------------------------------|--|
| Multiple Transports | Patented RDDA, NVMe-ready |
| Logical Volumes | Concatenated, RAID 0, RAID 1, RAID 10 |
| Multiple Drive Types | NVMe, NVMe, SATA, SAS |
| Failure Domains | Host, rack, and row aware |
| Management & Monitoring | |
| Interactive Interfaces | Web GUI and CLI commands |
| Automated Provisioning | RESTful API, Docker Persistent Volumes |
| Next Gen Data Center | |
| Flexible Topologies | Physically converged, disaggregated, or mixed |
| High Performance | Limitlessly scalable at near 100% linear performance |
| Scale-Out Architecture | Client, target, and management scale independently |
| Intelligent Clients | Clients utilize multiple hosts, drives, network-paths, and racks |
| Advanced Networking | Ethernet (RoCE v2), InfiniBand* |

About Excelerio

Excelerio enables enterprises and service providers to design scale-out storage infrastructures leveraging standard servers and high-performance flash storage. 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 a software-defined block storage solution that meets performance and scalability requirements of the largest web-scale and enterprise applications.

With Excelerio's NVMesh, customers can build distributed, high-performance server SAN for mixed application workloads. 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. The solution has been deployed for hyper-scale Industrial IoT services, machine learning applications and massive-scale simulation visualization.

Learn more about Excelerio NVMesh at <https://www.excelerio.com/>

Learn more about Intel storage solutions at www.intel.com/storage

Intel® Xeon® Scalable processors: intel.com/xeonscalable

Intel® Builders: <http://builders.intel.com>

Follow us on Twitter: [#IntelBuilders](https://twitter.com/IntelBuilders)



¹ Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Configurations: Performance claims obtained from data sheet, sequential read/write at 128k block size for NVMe and SATA, 64k for SAS. Intel® SSD DC P3700 Series 2 TB, SAS Ultrastar® SSD1600MM, Intel® SSD DC S3700 Series SATA 6 Gbps. Intel® Core™ i7-3770K processor @ 3.50 GHz, 8 GB of system memory, Windows® Server 2012, IOMeter. Random performance is collected with 4 workers each with 32 QD.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

© 2018 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon, are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

© 2018 Excelerio. Excelerio and NVMesh are trademarks of Excelerio in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.