

Excelero High-Performance Server SAN with Samsung NVMe SSDs

Excelero NVMesh[®], powered by Samsung NVMe SSD, provides software-defined block storage for scale-out applications.



Maximize scale-out block storage performance with Samsung NVMe and Excelero NVMesh

Summary

Excelero is revolutionizing the storage market with a software-based Server SAN solution NVMesh that provides high-performance block storage for scale-out applications. With Excelero's NVMesh, customers benefit from the performance and industry leading IOPS/\$ of Samsung NVMe SSDs, the convenience of centralized storage and the cost savings of standard server hardware. You get all these benefits while retaining local flash performance and latency characteristics. NVMesh was designed to meet the performance and scalability requirements of the largest web-scale and enterprise applications. The solution has been deployed by customers running hyper-scale Industrial IoT services, machine learning applications and massive-scale visualization of the Earth's oceans in simulated renderings.

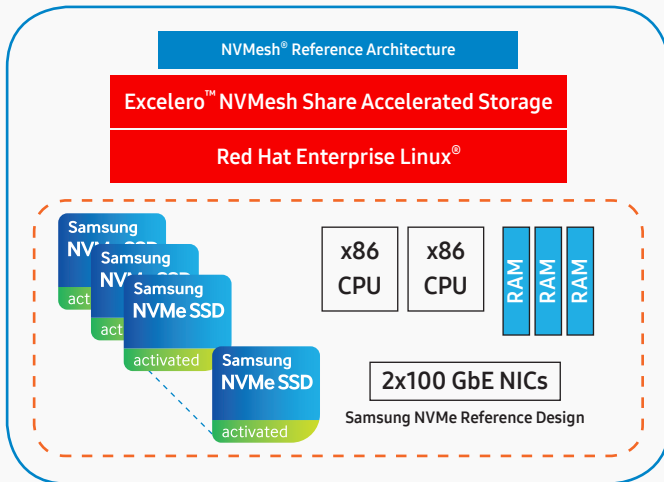


Figure 1: NVMesh/NVME Reference Architecture

The Excelero Reference Architecture which consists of the Excelero NVMesh running on the Samsung NVMe reference design can deliver close to 5 million random read 4K IOPs. With large blocks sizes (64KB), the Reference Platform can deliver 24GB/s of bandwidth. If that's not enough, multiple units can be seamlessly unified to reach virtually unlimited aggregate levels of bandwidth and IOPS. The Samsung NVMe Reference Design is engineered to provide perfectly-balanced storage nodes that include matching CPUs, networking, storage and PCIe connectivity to deploy large numbers of NVMe SSDs as a high-performance, yet low cost distributed block storage pool.

The result is a unified pool of high-speed NVMe storage, designed for web-scale deployments, high-performance enterprise applications or real-time analytics engines for IoT. Excelero's NVMesh volumes can also be shared by multiple hosts and/or applications.

Markets and Workloads

Excelero NVMesh with Samsung NVMe SSDs can be used for the most demanding high-performance, scale-out applications:

- Machine learning and Artificial Intelligence
- Real-time analytics
- Genomics and bioinformatics
- Scale-out web applications
- Enterprise databases
- Container Environments

Solution Benefits

- High performance per system enabled by NVMe SSD's
 - 4.9 million random read 4K IOPS
 - 4GB/s of Bandwidth
 - Very substantial high IOPS/\$
- Linear scalability with additional units
- Pure block storage, can be integrated with legacy applications or file & object infrastructures
- Open source hardware that can be customized to meet any deployment requirements.

With a Samsung PM953, the maximum capacity per system is 46 TB. With the next generation PM963 SSDs, the maximum capacity per system is 92 TB, and when using the high-endurance PM1725a, the maximum capacity per system is 153 TB. This reference platform is a dual socket Xeon®-based system with an Environmental Impact Assessment (EIA) compliant 2RU chassis. It uses 2x100Gb/s networking connectivity with Remote Direct Memory Access (RDMA). The Samsung NVMe Reference Design system is available through StackVelocity® (a business unit of Jabil Systems) as the Greyguard platform.

NVMesh Technology

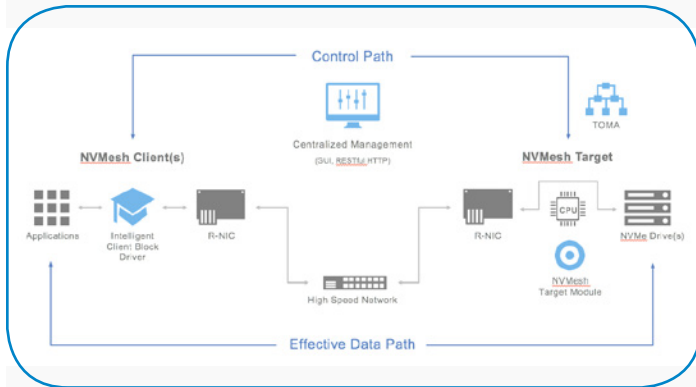
Excelero provides a software-based Server SAN solution that enables digital businesses to deploy high-performance block storage for scale-out applications. Excelero's NVMesh runs on standard servers with Samsung NVMe SSDs. NVMesh enables sharing NVMe SSDs on a network with equivalent performance and latencies to local direct attached SSDs, allowing applications to access a single pool of very high performance storage. The result is a high-performance block storage Server SAN that is a pre-validated, configured and fully supported solution. NVMesh features built-in logical volumes, redundancy, and centralized management.

Achieve the lowest \$/IOPS for NVMe-based storage at datacenter scale

NVMesh by Exceero is a high-performance, low-latency Software-Defined Storage (SDS) product. It provides remote, high speed, low latency storage facilities with in-server flash performance from standard off-the-shelf components. It allows efficient utilization of NVMe, NVmf and other SSD drives, potentially spread over many physical systems, treating them as a unified, redundant storage pool. In part, this is accomplished leveraging Exceero's patented Remote Direct Drive Access (RDDA) functionality.

NVMesh is comprised of four main software components:

- Centralized Management
- NVMesh Target Module
- Topology Manager (TOMA)
- Intelligent Client Block Driver



Samsung NVMe SSDs are increasingly being used as data storage media in computing, communication and multimedia devices, and offer superior reliability and value compared to traditional HDDs.

Advances in semiconductor flash memory have enabled the development of SSDs that are much larger in capacity compared to HDDs and can be used as direct replacements to them. SSDs also have proven to be highly cost-effective, due to their much lower power consumption and maintenance costs. As the world leader in semiconductor memory technology, Samsung revolutionized the storage industry by shifting planar NAND to a vertical structure. Samsung V-NAND technology features a unique design that stacks 48 layers on top of one another instead of trying to decrease the cells' pitch size. Samsung offers a comprehensive range of SSDs for deployment in a wide range of devices across virtually every industry segment.





Figure 3: Performance and Low Latency benefits of Samsung NVMe SSD's

Samsung NVMe Reference Design: The Samsung NVMe Reference Design system is a high-performance all-flash, scale-out storage server with up to 24 x 2.5-inch hot-pluggable Samsung advanced NVMe SSDs that provides extremely high capacity in a small footprint. It is based on PCIe Gen3 NVMe SSDs, offering the lowest latency in the industry with an optimized data path from the CPU to the SSDs. Each SSD slot provides power and cooling for up to 25 W per SSD to enable the support of current and future-generation, large-capacity SSDs, as well as SSDs with different endurance and performance levels.



Figure 4: Front Loaded and hot plug 24x NVMe SSDs

Build a performance tier that meets the scalability requirements of your applications

With exceptional balance, the Samsung NVMe Reference Design system allows performance to scale more linearly, without tending to be overprovisioned along any component. With Exceero NVMesh capabilities, enterprises can now unify Samsung enterprise NVMe SSD's across the data center as a single high-performance block storage pool reaching millions of IOPS.

Features & Capabilities

Combining blistering speeds with software-defined flexibility, NVMesh takes storage to its next generation. It transforms the performance, the economics and even the feasibility of multiple use cases in Web-scale IT, real time big data analytics, cloud infrastructure, and high performance databases.

NVMesh is a 100% software solution. Customers can choose servers, networks and NVMe media and combine them to tailor solutions with the latency, bandwidth and endurance parameters that fit their needs. They can also grow the system easily, simply and inexpensively.

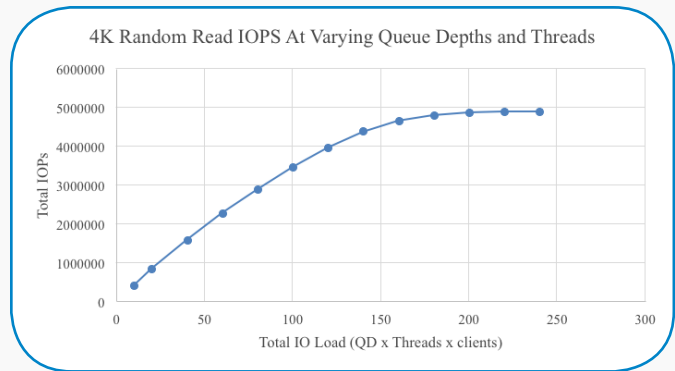
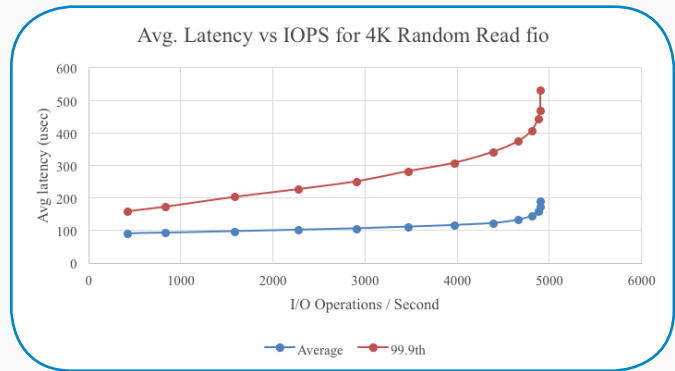
Some features of the combined Exceero NVMesh and Samsung NVMe Reference Design are:

- Pure block storage, can be integrated with legacy applications or file & object infrastructures
- High performance
 - 4.9 million random read 4K IOPS
 - 24GB/s of Bandwidth
- Reference Architecture is based on extensive testing jointly undertaken by Exceero and Samsung to characterize an optimized configuration
- Ability to use any standard servers
- NVMesh CPU offload technology allows the use of CPUs costing 10 times less than competitive solutions
- Striping and replication across nodes, enabling data durability, high availability and high performance
- Ability to add additional units at any time to expand existing volumes or grow the storage pool and support large capacity and scale out deployments.
- Upgrade clusters in phases—adding or replacing cards online—with minimal or no downtime
- Lower power consumption and higher reliability than similar-capacity HDD configurations
- Very substantial IOPS/\$

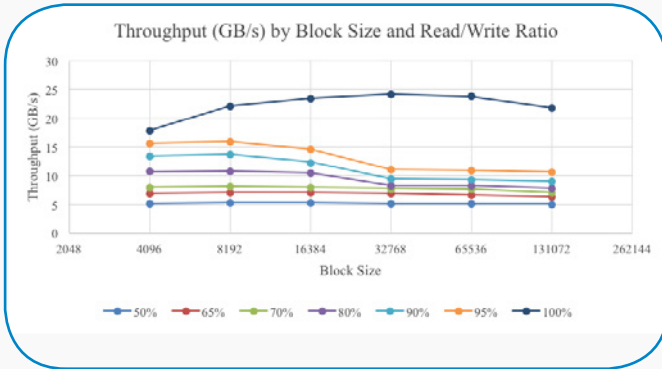
Performance

Equivalent performance to local storage

4.9M random read 4k IOPS at 150µs average latency, 99.9% under 450µs



Technical details



Technical details

Samsung and Excero performed extensive testing of a high-IOPS NVMe Storage cluster running over the Samsung NVMe Reference Design. Below is the Reference Architecture configuration, as well as the metrics of the combined solution.

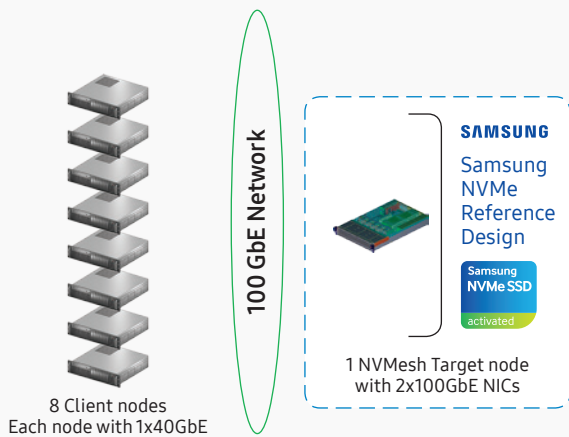


Figure 6: Samsung NVMe Reference Architecture Cluster

NVMe Reference Architecture

NVMe slots	<ul style="list-style-type: none"> 24 x 2.5 inch Samsung NVMe SSD slots Each slot supports up to 25 W per SSD Supports Samsung NVMe SSD's (PM963a and PM1725a)
CPU	2 x Intel® E5 v4 Series
Memory slots	Up to 512GB (Minimum 128GB)
Network	Mellanox® ConnectX®-4 EN: 2 x 100 GbE for network connectivity
Version of RHEL and NVMe	RHEL 7.3, NVMe 1.1.0, PM963a
Number of NVMe nodes	1 all-flash NVMe target node, 2 client nodes
Availability/redundancy	1+1 1200 W power supplies, 4+1 redundant fans
Remote accessibility	Dedicated 1GbE BMC (KVM/IP, IPMI)
Form factor	2U EIA-310-D, L 28", H 3.43", W 17.15", UL, CE, FCC, RoHS

Performance Highlights

IOPS (100% random read, IO size: 4 KB)	4.9M at 150µs average
IOPS (100% random write, IO size: 4 KB)	690K
Throughput (100% random read, IO size: 32 KB)	24.31GB/s

Legal and additional information

About Samsung Electronics Co., Ltd.

Samsung Electronics Co., Ltd. inspires the world and shapes the future with transformative ideas and technologies. The company is redefining the worlds of TVs, smartphones, wearable devices, tablets, cameras, digital appliances, printers, medical equipment, network systems, and semiconductor and LED solutions. For the latest news, please visit the Samsung Newsroom at news.samsung.com.

For more information

For more information about the Samsung NVMe Reference Design, please visit: <http://www.samsung.com/semiconductor/afard/>

About Excelero

Excelero 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 Excelero'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. For more information about Excelero and NVMesh, please visit: www.excelero.com.



Copyright © 2017 Samsung Electronics Co., Ltd. All rights reserved. Samsung and AutoCache are trademarks or registered trademarks of Samsung Electronics Co., Ltd. Specifications and designs are subject to change without notice. Nonmetric weights and measurements are approximate. All data were deemed correct at time of creation. Samsung is not liable for errors or omissions. All brand, product, service names and logos are trademarks and/or registered trademarks of their respective owners and are hereby recognized and acknowledged.

Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

The MariaDB trademark is wholly owned by MariaDB Corporation Ab and is a registered trademark in the United States of America and other countries.

Mellanox and ConnectX are registered trademarks of Mellanox Technologies, Ltd.

MySQL is a registered trademark of Oracle and/or its affiliates.

The OpenStack Word Mark and the OpenStack logos are trademarks of the OpenStack Foundation.

StackVelocity is a registered trademark of Jabil Circuit, Inc.

NVMesh is a registered trademark by Excelero.

Samsung Electronics Co., Ltd.
129 Samsung-ro,
Yeongtong-gu,
Suwon-si, Gyeonggi-do 16677,
Korea

www.samsung.com

2017-04