



# Delivering Fully Automated, Scale-Out NVMe<sup>®</sup>-oF Storage for AI & Advanced Analytics

High performance, continuously adaptable infrastructure for data-intensive applications on any OS or hypervisor

## EXECUTIVE SUMMARY

Organizations are investing in data and analytics to optimize their business operations, better engage with customers and gain a competitive edge, which in turn is driving the growth of artificial intelligence (AI), machine learning (ML), and advanced analytics. With a high demand for performance and low cost, these data-intensive applications most often run on bare-metal infrastructure, and now, NVMe flash.

Today, deploying bare-metal resources, expanding or reducing configurations, and gaining full utilization of those resources is challenged by complexity and inefficiency. The optimal solution will marry the flexibility and scale of networked storage with the performance and cost-efficiencies of local drives.

Together, DriveScale and Mellanox are making it possible to instantly provision NVMe storage over a 25G/100G Ethernet network to any host: bare-metal, Kubernetes, or hypervisor orchestrated.

## INFRASTRUCTURE CHALLENGES

NVMe is fast replacing SAS flash drives in the data center due to NVMe's enormous capacity, performance, and power advantages. However, the way in which NVMe storage is deployed impacts capacity utilization and performance results. NVMe flash drives today are as large as 15TB and growing ever larger. Putting NVMe drives in a server reduces flexibility and commonly results in underutilization. Any unused capacity is wasted since it can't be shared with other hosts. In addition, if storage requirements surpass what is local to the host, then there needs to be a way to extend the application to external storage.

External NVMe storage assigned to hosts over a network comes in two forms: (1) SAN arrays and file storage that incorporate proprietary controllers and data management software, and (2) commodity storage. Both storage types are capable of NVMe-oF,

## HIGHLIGHTS

- Automated provisioning of networked, highly adaptable storage to servers without using any host resources
- Data-intensive applications on any OS or hypervisor can now take advantage of scale-out NVMe storage
- Server off-load for storage orchestration, encryption, RAID and other services to preserve application host performance
- Ability to instantly attach/detach data sets and storage, and replace failed components in seconds
- Local-drive performance at scale and throughput that scales with demand
- Scale-out architecture in any size increments

whether NVMe/RDMA or NVMe/TCP. In the case of traditional enterprise SAN or file storage, the cost to scale is high, and performance impedance occurs as unneeded proprietary software gets in the way of I/O. And more importantly, data-intensive applications such as NoSQL, distributed SQL, Spark, Kafka and machine learning do not need costly additional data management services such as copy, replication, deduplication and backup that are embedded in low-performance traditional arrays.

With commodity NVMe systems widely available, scale-out, cost-efficient solutions can better meet the performance, and low-cost demands of data-intensive, modern applications. The DriveScale Composable Platform is designed to provide the automation, flexibility and adaptability that makes commodity storage functional. With the programmable Mellanox BlueField® SmartNIC, users gain the advantage of server offload for storage management and services in addition to security and network services. In addition, DriveScale's server agent runs on the BlueField SmartNIC enabling DriveScale to support application hosts running any operating system or hypervisor.

## ELASTIC BARE METAL CLOUDS WITH COMPOSABLE INFRASTRUCTURE

Bare metal cloud infrastructure needs to be adaptable and flexible to meet the needs of multiple applications and groups that share the infrastructure. With DriveScale Composable Platform and Mellanox BlueField SmartNIC, users get high performance, automated and programmable resources that scale efficiently. The DriveScale server agent running on the BlueField SmartNIC offloads storage orchestration and services from the host ensuring applications have access to all of the host's processing capacity.

DriveScale automates the NVMe data fabric setup – whether NVMe/RDMA, NVMe/TCP, or NVMe/iSCSI. And NVMe flash drives or slices of drives appear as local drives to the server or hypervisor with local drive performance. With DriveScale and Mellanox, users gain the cost-efficiencies of commodity storage with the flexibility of networked storage over a standard 25G/100G Ethernet network.

## DRIVESCALE, MELLANOX ENABLE SCALE-OUT, HIGH-PERFORMANCE INFRASTRUCTURE

DriveScale Composable Platform when used with Mellanox BlueField and Spectrum Switches enables an end-to-end solution that automates the discovery of servers, storage, and data fabric capabilities, and configures them into adaptable resource clusters. DriveScale works with Mellanox to detect data fabric connectivity options offering choice and reduced complexity in preparing the data fabric. Users choose the instances of compute, and the instances of storage needed for a workload and DriveScale's software platform automatically sets up the highest performing fabric, mounts the devices, and creates a cluster ready for workload deployment.

## DRIVESCALE AND MELLANOX FOR AI

For machine learning and deep learning workloads, DriveScale provides unique capabilities to ensure full utilization of GPUs and to eliminate wasted dataset copying. Rather than copying datasets into the data pipeline, users can attach and detach datasets to GPUs in just 1 second, eliminating wait times that plague GPU systems. In addition, each stage of the data pipeline has different requirements for performance, latency, and reads. With DriveScale, the optimal configuration for each stage can be set up, adapted, and torn down in seconds as needed.

Moving your static AI infrastructure to a programmable, adaptable, and API-driven solution will save on operational and resource inefficiencies. The ability to deploy physical resources on the fly makes it possible to match the demands of each stage of the data pipeline with the right resource configuration. And with DriveScale, all available bandwidth to the storage is accessed, providing industry-leading ingest rates.

## STORAGE OFFLOAD AND HIGH-PERFORMANCE FABRIC

Mellanox BlueField SmartNICs support NVMe-oF, a high-performance storage protocol designed to take advantage of faster flash storage over RDMA or TCP. For NVMe/RDMA, Mellanox provides RoCE, a network protocol that allows remote direct memory access over Ethernet, offloading the data transfer functions to the adapter to bypass the CPU. These technologies provide for a more efficient and faster way to move data while lowering latencies and CPU utilization.

With more powerful CPUs and GPUs and faster storage required in data analytics, the switch is at the heart of the network, and must be able to keep pace with fast and intensive data movement. This is accomplished by moving data at low latency between CPUs, GPUs, memory, and storage. Mellanox Spectrum™ Ethernet switches provide 25/100GbE line-rate performance and consistent low latency with zero packet loss. Spectrum is also the only RoCE-ready switch that can deploy RoCE effortlessly, offer end-to-end automatic RoCE acceleration and real-time RoCE visibility for easy troubleshooting.

## CONCLUSION

DriveScale and Mellanox transform industry-standard compute nodes, GPU nodes, and storage systems into elastic resource instances enabling users to deploy IT resources on-demand and expand, reduce or replace resources as workloads require. Attach and detach datasets to GPU compute nodes instantly to fully utilize GPU resources.

Mellanox high-performance, low-latency fabric combined with accelerators and offloads ensure the proper fabric to provide efficient data for analytics and ensure future scalability. By deploying DriveScale on a Mellanox infrastructure, IT organizations can reduce cost and complexity while driving up CPU efficiency and GPU utilization. Enterprises can achieve hyperscale performance at a fraction of the cost of public cloud solutions and without the need for a highly trained staff.



## WANT TO LEARN MORE?

DriveScale Composable Platform:

<https://drivescale.com/composable-platform/>

Mellanox BlueField SmartNIC:

<https://www.mellanox.com/products/bluefield-overview>

Mellanox Spectrum Ethernet Switches:

[https://www.mellanox.com/page/products\\_dyn?product\\_family=218&mtag=spectrum\\_ic](https://www.mellanox.com/page/products_dyn?product_family=218&mtag=spectrum_ic)

### About DriveScale

We are a team of entrepreneurs and experts in diverse disciplines who are united by a common quest: to re-imagine static compute infrastructure as adaptable and programmable, making every data center a composable, agile, scale-out cloud. This gives any company the ability to enjoy the same speed, flexibility, and cost-efficiency that the Cloud Giants create for themselves. With the growth of data-driven and distributed applications, enterprise and cloud companies will benefit from the rapid deployment, flexible operations and high availability previously only found in the largest data centers in the world. This is our vision for your infrastructure.

For more information, visit [www.drivescale.com](http://www.drivescale.com)

### About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables and silicon that accelerate application runtime and maximize business results for a wide range of markets including high-performance computing, enterprise data centers, Web 2.0, cloud, storage and financial services.

To find out more, visit our website: [www.mellanox.com](http://www.mellanox.com)



350 Oakmead Parkway, Suite 100  
Sunnyvale, CA 94085  
Tel: 408-970-3400 • Fax: 408-970-3403  
[www.mellanox.com](http://www.mellanox.com)