

DriveScale NVMe Adapter

The DriveScale Composable Platform automates the creation of high-scale, high-performance clusters of commodity, heterogeneous disk-lite server, storage and network fabric elements enabling IT to compose optimized cluster configurations for each workload from an intelligent software platform. This enables the quick provisioning of compute and storage in a cluster, the flexible re-allocation of compute and storage capacity, and the ability to assign additional compute or storage resources on demand. Compute can be upgraded separately from storage to optimize spend and take advantage of processor performance updates without having to upgrade disk at the same time. With DriveScale, resources can be more fully utilized eliminating wasted spend.

DriveScale eliminates the constraints and high costs of direct-attached storage (DAS) servers and appliances by disaggregating compute and storage elements and connecting them to a standard Ethernet network. With high network interface speeds and network bandwidth available today, it is now possible to attach storage directly to the network and provide similar performance to an integrated server chassis. DriveScale and third-parties provide both SAS to Ethernet and NVMe to Ethernet adapters to attach storage to an Ethernet network. Here we describe the DriveScale NVMe Adapter.

DriveScale NVMe Adapter

For high performance, latency-sensitive workloads, flash is the optimal solution. The DriveScale NVMe Adapter enables users to connect NVMe solid-state storage (SSD) to an Ethernet network. It also makes it possible to better utilize flash resources by enabling slicing – the ability to mount portions of an SSD to individual servers. By combining NVMe flash drives with the DriveScale Composable Platform, customers can optimize the configuration of compute to flash capacity in their cluster for applications such as NoSQL, Cassandra, Vertica, Apache Spark, and Kubernetes.

Core to the design of the adapter is a high availability, high performance, low-latency, and secure architecture. The high availability design increases accessibility of your data and enables quick resolution of disk and compute node failures through a software interface. High I/O performance at scale is achieved through a low-latency design and software optimization. The DriveScale NVMe Adapter Software provides an inventory of storage resources and individual drive intelligence to the DriveScale Composable Platform. The software runs on standard Linux servers and exports the iSCSI or NVMe drives from a shared pool to compute nodes as local SCSI or NVMe block devices.

NVMe Adapter Features

- **NVMe and NVMeOF:** DriveScale creates a software stack that bridges NVMe to Ethernet and also supports NVMeOF with the RoCEv2 protocol enabling seamless connectivity of NVMe flash to the Ethernet network. Putting flash drives on the network via Ethernet Bunch of Flash (eBOFs) enables an optimized number of drives to be attached to servers eliminating over provisioning, lowering costs and improving utilization.
- **NVMe SSD slicing:** Solid-state drives provide a very high rate of I/O per second (IOPS) which makes them suitable for sharing across servers. The DriveScale Adapter Software enables slicing of an NVMe SSD into portions that can be used to serve many different servers for data that requires high throughput or low latency access. Slicing also enables a cost-effective way of utilizing an entire SSD.
- **High performance:** The Adapter Software supports DriveScale's near line-rate network speeds and high I/O performance delivering direct-attached drive performance for drives attached via the DriveScale Composable Platform.



DriveScale

DriveScale NVMe Adapter

- Security:** The Adapter Software enables authentication of servers and drives when directed by the DriveScale Composer. Drives attached to a server cannot be accessed by other servers.
- Inventory management and monitoring:** The Adapter Software monitors and reports the full SAS/PCIe connection tree including bandwidth, multiple paths, and shared connections. In addition, it reports network connectivity and retrieves and reports SMART and temperature information on the drive for failure detection.
- Load balancing:** The load balancing capability provides continuous load balancing capabilities between adapter modules, network path and interfaces for data transfer based on the current data load, interface link status, and congestion.

DriveScale NVMe Adapter Chassis

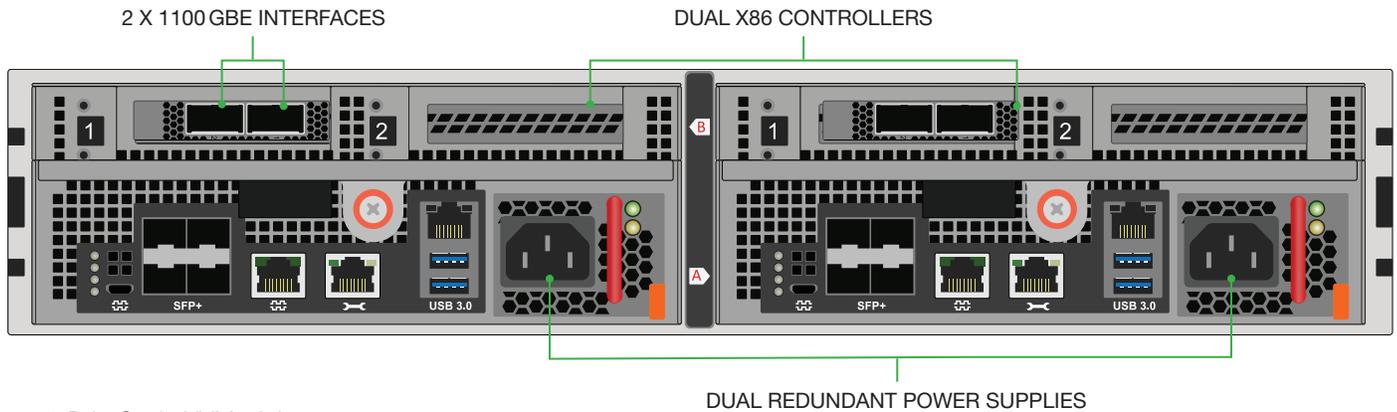


Figure 1: DriveScale NVMe Adapter

Features	Description
Form factor	2RU chassis
Number of controllers	2
Onboard 1GbE Management connectivity per controller	1
Mellanox ConnectX®-4 adapter cards per controller	1
Onboard 100GbE ports per controller	2
Dual redundant power supplies per adapter	Yes
Hot swappable	Yes
Number of 2.5" NVMe SSD slots	24
Support for third-party single and dual-port NVMe SSDs	Yes
CPU	Intel Skylake D
Supported drives	Intel, Western Digital
Protocols supported	iSCSI and NVMe over Fabrics (RoCE v2)
Dimensions	44.72 x 74 x 8.59 cm (17.61 x 29.13 x 3.39 in)
Weight	32 Kg
Input power	100-264V 50Hz/60Hz AC power at 1100W, IEC320-C14 power receptacle
Operating temperature	5°C to 40°C (41°F to 104°F) sea level to 3000 feet; derate 1°C (1.8°F) for every 1000 feet of elevation
Operating humidity	8% to 85% non-condensing
Regulatory approvals	RoHS, VCCI-A, CE, CAN ICES, UL