Accelerate Your AI Outcomes

Supercharge Workloads with DDN's High-Performance Data Platform

COMMON USE CASES

The flexibility and capabilities of the SFA NVMe platforms lend themselves to a wide variety of applications:

> Artificial Intelligence

> Analytics

> Deep Learning

> Content Distribution

> High IOPs Telemetry

DDN Storage Fusion Architecture® (SFA®) data platforms are purpose-built to deliver scalable flash performance and capacity to flexibly meet your changing business demands.

The SFA200X2 & SFA400X2 platforms devlier 100% NVMe storage with multiple high-speed connectivity options. With throughput up to 115 GB/s, the 2U SFA NVMe-only/Hybrid models are among the fastest data platforms in the industry, capable of supporting up to 24 NVMe SSDs in a compact form factor. Built with 3rd Generation Intel® Xeon® Scalable processors, this extreme level of performance density makes the SFA NVMe platforms ideal for data centers with limited space that require robust high-performance flash in a scalable, out-of-the-box architecture. Start with a single enclosure and scale limitlessly to meet file system or block storage requirements. Hybrid configurations support up to 24 NVMe devices, plus an additional 120 QLC or 900 SAS drives (spinning disk or SSD).

BLINDING PERFORMANCE THAT SCALES ON DEMAND

Whether you're accelerating analytics workloads, reducing latency for tough NoSQL databases, or launching a deep learning project with modest training sets, the SFA NVMe platforms are ideal as cost-effective building blocks. Designed to maximize the value of your investment, the platform includes an internal switch fabric with up to 288 lanes of PCIe Gen 4, along with up to 64 lanes for client connectivity. The SFA400NVX2 model can handle up to 16x x4 ports (64 lanes) of SAS-4 connectivity, offering a seamless hybrid expansion capability.

SMALL, YET POWERFUL FLEXIBILITY

SFA NVMe platforms are available as block storage appliances and as integrated high-performance file appliances. DDN's EXAScaler® file system extends the SFA NVMe storage building block model to scale-out the parallel file namespace with maximum efficiency. Built, deployed, and supported by experts in data-intensive workloads, these appliances deliver optimal application and workload performance by combining the advantages of DDN's SFAOS with the industry-leading EXAScaler® parallel file system.

TECHNICAL SPECIFICATIONS

	SFA200X2	SFA400X2	AI400X2®	SFA400X2T
SYSTEM FEATURES	Active/Active storage controllers DeClustered RAID (DCR) supports erasure coding schemas: RAID 6 8+2, 4+2; RAID 5 8+1, 4+1; RAID 1 1+1		
PERFORMANCE	Sequential read performance up to 48GB/s Sequential write performance up to 38GB/s Up to 1.5M IOPs	Sequential read performance up to 90GB/s Sequential write performance up to 65GB/s Up to 3M IOPs	Sequential read performance up to 115 GB/s Sequential write performance up to 75 GB/s Up to 3M IOPs
CPUS PER APPLIANCE	2x 3rd Gen Intel® Xeon® Scalable processors	4x 3rd Gen Intel® Xeon® Scalable processors	4x 3rd Gen Intel® Xeon® Scalable Processors
CONTROLLER HOST PORTS PER APPLIANCE	4x HDR/HDR100 or 200/100GbE	8x HDR/HDR100 or 200/100GbE	16x HDR/HDR100 or 200/100 GbE
DRIVE SUPPORT	24 x 2.5″ dual port, hot-swappable NVMe	24 x .25″ dual port, hot-swappable NVMe SAS-4 option: Up to 120 QLC SSD	24 x 2.5″ dual port, hot-swappable NVMe
STANDARD SOFTWARE FEATURES	LUN mapping and masking, intelligent write striping, read QoS, port zoning detection, data integrity check/correction, interface options (SSH to CLI, web-based GUI, Python API), state change messages (via e-mail, SNMP trap and syslog).		
AVAILABLE FILE	EXAScaler ES200NVX2	EXAScaler ES400NVX2, AI400X2	EXAScaler ES400NVX2T

ABOUT DDN

DDN is the world's leading data intelligence company that provides an unfair advantage to over 11,000 customers focused on unlocking real-time AI insights. The DDN Data Intelligence Platform supercharges more than 500,000 GPUs worldwide across a broad range of use cases, including autonomous driving, financial services, healthcare, research, and academia. Manage complex data, enhance performance, deliver cost savings, increase security and accelerate your AI workloads at-scale from edge to core to cloud.