

DDN A³I[®] STORAGE PLATFORMS

Providing unmatched storage performance and flexibility for Artificial Intelligence and Deep Learning

20x Scaling Efficiency

True parallel architecture performance, efficiency, GPU utilization, and storage capacity at any scale.

20x DL Acceleration

Delivers full, real-time acceleration for all workloads concurrently and continuously.

Proven at-Scale

Powers the largest NVIDIA DGX SuperPODs, and most extensive deep learning programs globally.

Rock Solid

Data protection, integrity, declustering, redundancy ideal for mission critical applications

Effortless Deployment

Fully-integrated and optimized for AI workloads and GPU enabled solutions.

DDN A³I solutions (Accelerated, Any-Scale AI) break new ground for AI, Analytics and HPC. Engineered from the ground-up for the AI-enabled data center, DDN's A³I solutions accelerate AI applications and streamline data-intensive workflows using the DDN shared parallel architecture. The DDN AI200X/AI400X all-NVME appliances and AI7990X capacity appliance provide unmatched flexibility for your organization's AI needs. DDN A³I solutions are easy to deploy and manage, highly scalable in both performance and capacity, and represent a highly efficient and resilient platform for all of your current and future AI requirements.

Fully Integrated, GPU-Optimized Data Platforms

Easy to deploy, A³I solutions are turn-key, pre-configured, and provide the most capable scale-out platform for capacity and performance. DDN A³I solutions ensure full saturation of GPUs for maximum productivity at any scale.

Full GPU Saturation

Fully optimized for all types of workloads and data, A³I solutions ensure full GPU resource utilization even with distributed applications running on multiple computing servers. Performance testing on the DDN A³I architecture has been conducted with all widely-used DL frameworks, and consistently demonstrates higher performance and faster completion times. Using the A³I parallel client, containerized applications can engage the full capabilities of the data infrastructure and servers achieve full GPU saturation for DL workloads. A³I solutions fully integrate NVIDIA GPUDirect™ Storage for most-efficient data access from GPUs.

Capacity-Efficient AI Storage

A³I solutions provide flexible capacity expansion options, with up to 256TB of scale-out NVMe capacity per AI200X/AI400X appliances, or 4PB of hybrid storage in the AI7990X.

Highest Resiliency, Reliability, and Security at Scale

Engineered to provide the highest data availability and maximum system uptime, all A³I hardware and software components are integrated as a fully redundant system.


Unified Namespace

A³I solutions allow for consolidation of hot training data and warm expanding data libraries into a single platform, providing easy data access from a unified interface. Additional data services are available, including S3 interface for seamless file and object data access.


Multi-Tenancy and Quota Support

A³I solutions can be secured on a per-tenant basis that ensure users and applications can only access the data that they're entitled to. Advanced quota tools provide easy data management.

TECHNICAL SPECIFICATIONS



AI200X™



AI400X™



AI7990X™

System Features

High performance GPU-optimized parallel file system

Sequential read performance
up to 24GB/s

Sequential write performance
up to 20GB/s

Up to 1.5M IOPs per appliance

High performance GPU-optimized parallel file system

Sequential read performance
up to 48GB/s

Sequential write performance
up to 34GB/s

Up to 3M IOPs per appliance

High performance GPU-optimized parallel file system

Sequential read performance
up to 24GB/s

Sequential write performance
up to 20GB/s

Up to 800K IOPs per appliance

Controller Host Ports per Appliance

4 x EDR/HDR100* InfiniBand or 100 GbE

8 x EDR/HDR100* InfiniBand or 100 GbE

4 x EDR/HDR100* InfiniBand or 100 GbE

Drive Support

2.5" dual port NVMe drives

32TB, 64TB, 128TB, 256TB useable capacity
configurations

2.5" dual port NVMe drives

32TB, 64TB, 128TB, 256TB useable capacity
configurations

Enterprise-grade HDDs (data) and
SSDs (metadata)

1PB, 2PB, 4 PB useable capacity configurations

Standard Software Features

High performance parallel file system, 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).

Safety

Agency Certifications UL, cUL, CE, FC

Agency Certifications UL, cUL, CE, FC

Agency Certifications UL, cUL, CE, FC

Product Specifications Subject to Change. For physical and environmental attributes, see associated SFA Data Sheets *X-variant is HDR100 (releases starting Dec 2019), non-X remains EDR

©DataDirect Networks. All Rights Reserved. DataDirect Networks, the DataDirect Networks logo, DDN, AI, AI200X, AI400X, and AI7990X are trademarks of DataDirect Networks. Other names and brands may be claimed as the property of others.

v7 (4720)