



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.

Rock Solid and Proven at Scale

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[®] Storage Platforms

DDN's A³I (Accelerated, Any-Scale AI) solutions break new ground for Artificial Intelligence (AI) and Deep Learning (DL), providing unmatched flexibility for your organization's AI needs. Engineered from the ground-up for the AI-enabled data center, A³I solutions are optimized for ingest, training, data transformations, replication, metadata and small data transfers. DDN offers flexibility in platform choice with the all-flash NVMe AI200X/AI400X or AI7990X, a hybrid flash and hard drive storage platform which leverages parallel access to flash and deeply expandable HDD storage. The AI200X/AI400X and AI7990X support a scale-out model with solutions starting at a few TBs yet scalable to 10s of PBs.

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.

Full GPU Saturation

Fully optimized for all types of I/O patterns and data layouts, A³I solutions deliver data to applications, ensuring 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 (TensorFlow, Horovod, Torch, PyTorch, NVIDIA[®] TensorRT[™], Caffe, Caffe2, CNTK, MXNET and Theano). Using the A³I intelligent client, containerized applications can engage the full capabilities of the data infrastructure, and that the AI servers achieves full GPU saturation consistently for DL workloads.

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.

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 controls provide easy management of file system consumption at the user, group, and project level.

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 30GB/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

2.5" dual port NVMe drives

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

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

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

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 UI, cUI, CE, FC

Agency Certifications UI, cUI, CE, FC

Agency Certifications UI, cUI, 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

About DDN

DataDirect Networks (DDN) is the world's leading big data storage supplier to data-intensive, global organizations. DDN has designed, developed, deployed, and optimized systems, software, and solutions that enable enterprises, service providers, research facilities, and government agencies to generate more value and to accelerate time to insight from their data and information, on premise and in the cloud.

©DataDirect Networks. All Rights Reserved. DataDirect Networks, the DataDirect Networks logo, DDN, A1, AI200X, AI400X, and AI7990X are trademarks of DataDirect Networks. Other names and brands may be claimed as the property of others.

v6 (4/20)