



# DDN A<sup>3</sup>I<sup>®</sup> Storage Platforms

DDN's A<sup>3</sup>I (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, A<sup>3</sup>I 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 AI200X2/AI400X2 or AI7990X, a hybrid flash and hard drive storage platform which leverages parallel access to flash and deeply expandable HDD storage. The AI200X2/AI400X2 and AI7990X support a scale-out model with solutions starting at a few TBs yet scalable to 10s of PBs.

## 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.

### ▶ FULLY INTEGRATED, GPU-OPTIMIZED DATA PLATFORMS

Easy to deploy, A<sup>3</sup>I 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, A<sup>3</sup>I solutions deliver data to applications, ensuring full GPU resource utilization even with distributed applications running on multiple computing servers. Performance testing on the DDN A<sup>3</sup>I architecture has been conducted with all widely-used DL frameworks (TensorFlow, Horovod, Torch, PyTorch, NVIDIA<sup>®</sup> TensorRT<sup>™</sup>, Caffe, Caffe2, CNTK, MXNET and Theano). Using the A<sup>3</sup>I 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

A<sup>3</sup>I solutions provide flexible capacity expansion options, with up to 256TB of scale-out NVMe capacity per AI200X2/AI400X2 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 A<sup>3</sup>I hardware and software components are integrated as a fully redundant system.

### ▶ UNIFIED NAMESPACE

A<sup>3</sup>I 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

A<sup>3</sup>I 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



**AI200X2™**



**AI400X2™**



**AI7990X™**

### SYSTEM FEATURES

High performance GPU-optimized parallel file system

Sequential read performance up to 48GB/s

Sequential write performance up to 38GB/s

Up to 1.5M IOPs per appliance

High performance GPU-optimized parallel file system

Sequential read performance up to 90GB/s

Sequential write performance up to 64GB/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

4x HDR/HDR100 or 100/200GbE

8x HDR/HDR100 or 100/200GbE

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

### 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, Inc. All Rights Reserved. DataDirect Networks, EXAScaler, ES200NVX, ES400NVX, ES7990X, and ES18KX are trademarks of DataDirect Networks. \*Other Names and Brands May Be Claimed as the Property of Others.