DDN A³I® Solutions

DDN A³I solutions (Accelerated, Any-Scale AI) break new ground for artificial intelligence (AI) and big data analytics. Engineered from the ground up for the AI-enabled data center, DDN's A³I solutions accelerate applications and streamline data-intensive workflows using the DDN shared parallel architecture. The DDN AI400X™ all-NVME appliance is the building block that provides unmatched flexibility for your organization's AI needs.

Optimize your entire Enterprise AI environment

33x 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.

Effortless Deployment

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

From the AI Experts

DDN is recognized global leader in at-scale data. Unrivaled experience and expertise with strong-customer focus and commitment.

Simplified AI Infrastructure, Proven at Every Scale

Easy to deploy, DDN A³I solutions are tightly integrated with NVIDIA® DGX™ systems for turnkey AI infrastructure and provide the most capable scale-out platform for capacity and performance.

The simplicity of the solution makes it appropriate for AI projects of all sizes. From individual DGX A100s and DGX POD™ to the largest DGX SuperPOD™, DDN's parallel architecture maximizes the utilization of the entire AI infrastructure for true application acceleration.

Efficient AI Storage for Maximum Application Performance

Our AI solutions deliver high throughput, low latency, and massive concurrency with a shared parallel architecture. The DDN AI400X packages all the capabilities needed for unified AI storage into an efficient 2U appliance that delivers 50 GB/s throughput. Fully optimized for all workloads and data types, DDN A³I solutions ensure full GPU resource utilization for applications running on multiple computing servers. DDN has fully integrated NVIDIA's GPUDirect™ Storage for most-efficient data access from GPUs and we continue to collaborate with NVIDIA on developments like NVIDIA BlueField®.

Predictable Scaling

Pre-configured appliances with defined capacity, performance and capability are simple building blocks that scale with AI project success. Supporting hybrid expansion using cost effective HDDs, customers can scale-out with all flash performance or grow a single system over 6PB for longer term retention.

Intelligent Automated Data Management

DDN 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. Data is managed automatically between flash and capacity disk storage without administrator intervention. Additional data services including NFS and S3 interfaces for easy ingest and sharing from one platform. Secure multi-tenancy and advanced quota tools make AI as a service delivery simple for enterprise-wide or service providers.

Technical Specifications

AI400X™

System Features

High performance GPU-optimized parallel file system

Sequential read performance
up to 50GB/s

Sequential write performance
up to 34GB/s

Up to 3M IOPs per appliance

Dual, hot swappable power supplies

Controller Host Ports per Appliance

8 x EDR/HDR100* InfiniBand or 100 GbE

Drive Support

2.5” dual port NVMe drives

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

Hybrid deployment with Enterprise-grade HDDs

Up to 6.7PB usable capacity configurations available

Software Features

High performance parallel file system, automated data tiering, multi-tenant security, encryption, 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

Physical and Environmental Attributes

Dimensions

Height: 2RU rack mount 3.5” (89 mm)

Width: 19” rack (482.6 mm)

Depth: 33.5” (850 mm) without bezel

Power/Cooling

Input Voltage: 200-240V 50/60 Hz
Nominal Power: 675 W (empty); 1,275 W (max)

Nominal Heat

2,303 BTU/hr (empty); 4,350 BTU/hr (max)

Operating Environment

Op Temp, Sea Level: 10-35 degC; Op Temp, 3000m: 10-28 degC; 20%-80% humidity range

Weights

80 lbs/36 kg (empty); 90 lbs/41 kg (max)

Product Specifications Subject to Change Without Notification.

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.

v10 (421-KK)