Optimizing EXAScaler for AI workloads and AI infrastructure.


DDN is the World's Largest Privately Held Storage Company
Market Leader With Global Presence
10,000
Customers
1,000
Employees
150+
Patents
20 Years
of Industry Leadership
10
Technology Centers Around the World

Powering the World's Data Intensive Workflows In All Industries

AI & Analytics

Optimized Turnkey End-to-End AI Solutions Engineered for All Your Use Cases at Any Scale

Web and Cloud

Enhancing Your On Prem and Private Clouds With the Most Versatile Public Cloud Storage

Enterprise

Combining Reliability and Business Continuity with Enterprise Analytics and Management Tools

Enterprise at Scale

Bringing Massive Scale and True Flexibility to the World of Enterprise Data Storage

Govt/Academia

20 Years of Delivering the Most Powerful At Scale Data Solutions For All Your HPC Infrastructures

EXAScaler Platforms and Specifications

NVMe All-Flash	All NVMe/Hybrid	Mid Range Hybrid	EXAScale Class
		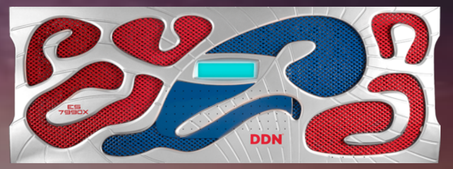	
ES200NVX	**ES400NVX**	**ES7990X**	**ES18KX**
24GB/s **1.5M IOP/s**	**50GB/s** **3M IOP/s**	**24GB/s** **800K IOP/s**	**76GB/s** **3M IOP/s**
24 NVME Slots	24 NVME Slots Up to 360 SAS Slots	Up to 450 SAS Slots	48 NVMe Slots Up to 1872 SAS Slots
EDR/HDR100 IB (4) Or 40/100GbE (4)	EDR/HDR100 IB (8) Or 40/100GbE (8)	EDR/HDR100 IB (4) Or 40/100GbE (4)	EDR/HDR100 IB (16) Or 40/100GbE (16)

Performance Simplified with EXAScaler Appliances

CLIENT → HCA/NIC → SWITCH → HCA/NIC → SERVER → FILESYSTEM → HBA → SAN SWITCH → HBA → STORAGE

TRADITIONAL

IO PATHS

Consolidated, Simplifed Componentry with no compromise on performance

EMBEDDED

CLIENT → HCA/NIC → SWITCH

The Fastest Data Solution for Trading

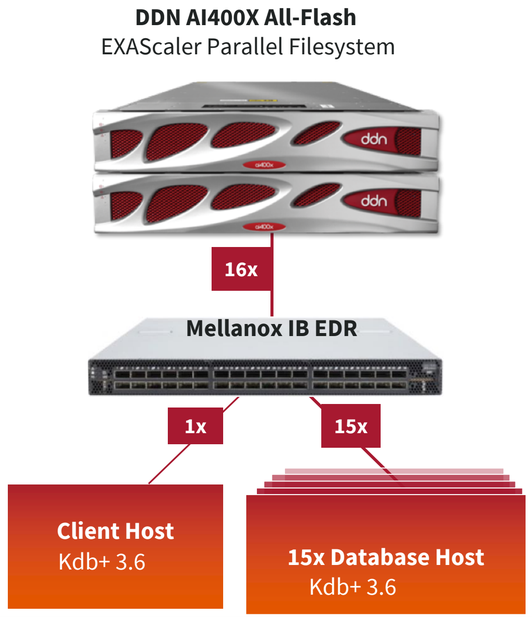

DDN Tested Platform

- 72% less footprint
- 1/3 of the power consumption
- 1/4 of the Operation Cost

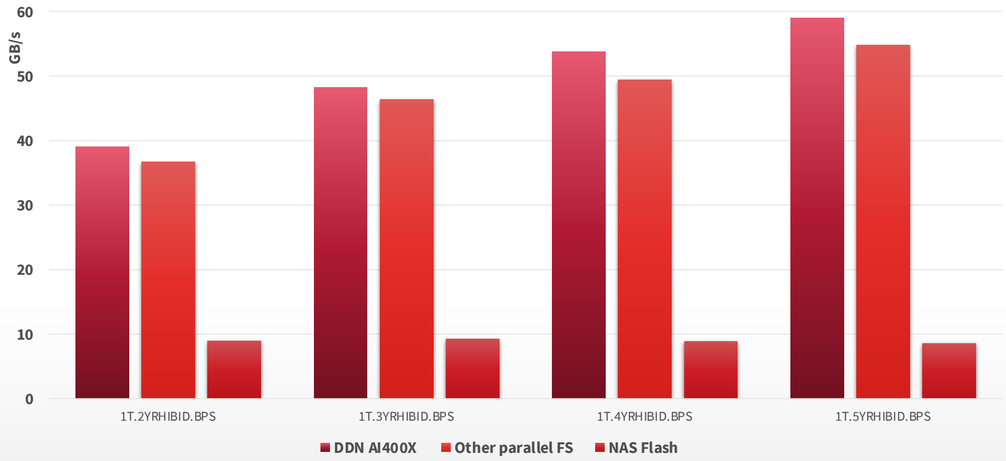

www.STACresearch.com/KDB200915

The Fastest Cloud Data Platform

Available on GCP, AZURE, and AWS*

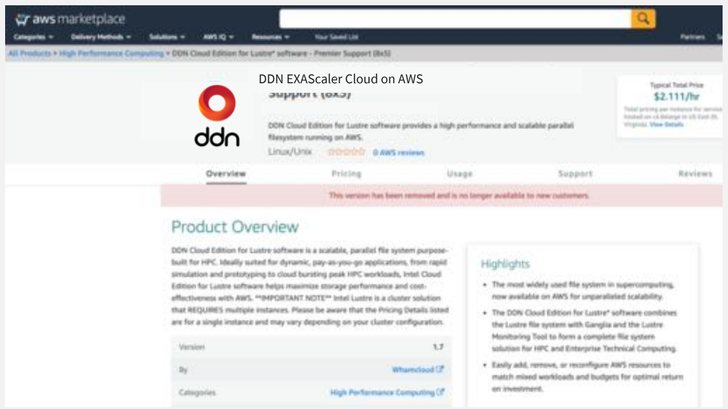

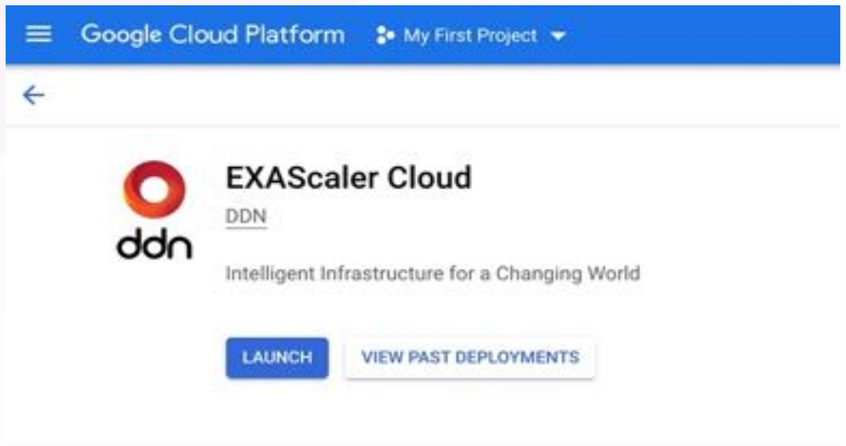

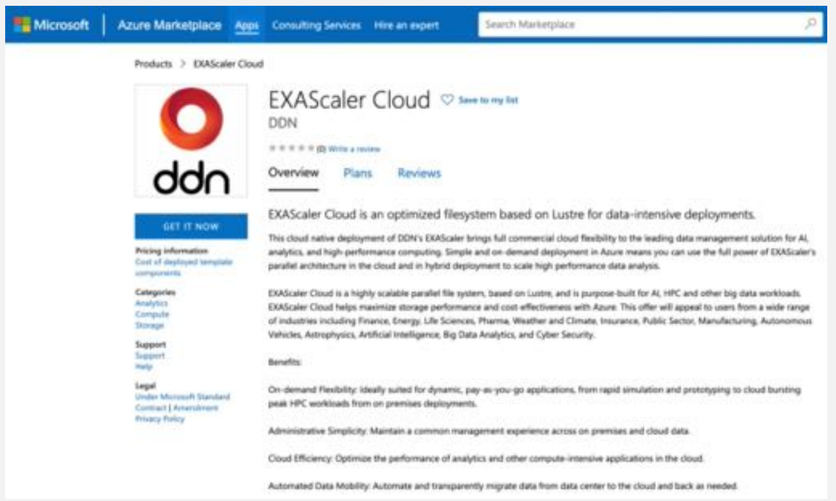

Cloud Performance

400TB EXAScaler filesystem with 32 clients shows aggregated throughput of over 50GB/s

Node	Count	Type	Disk	Size
MDS	1	Skylake n1-stan	dPaDr-dS-S4D	1TB
OSS	40	Skylake n1-stan	dPaDr-dS-T1D6	10TB
Client	40	Skylake n1-standard-8		

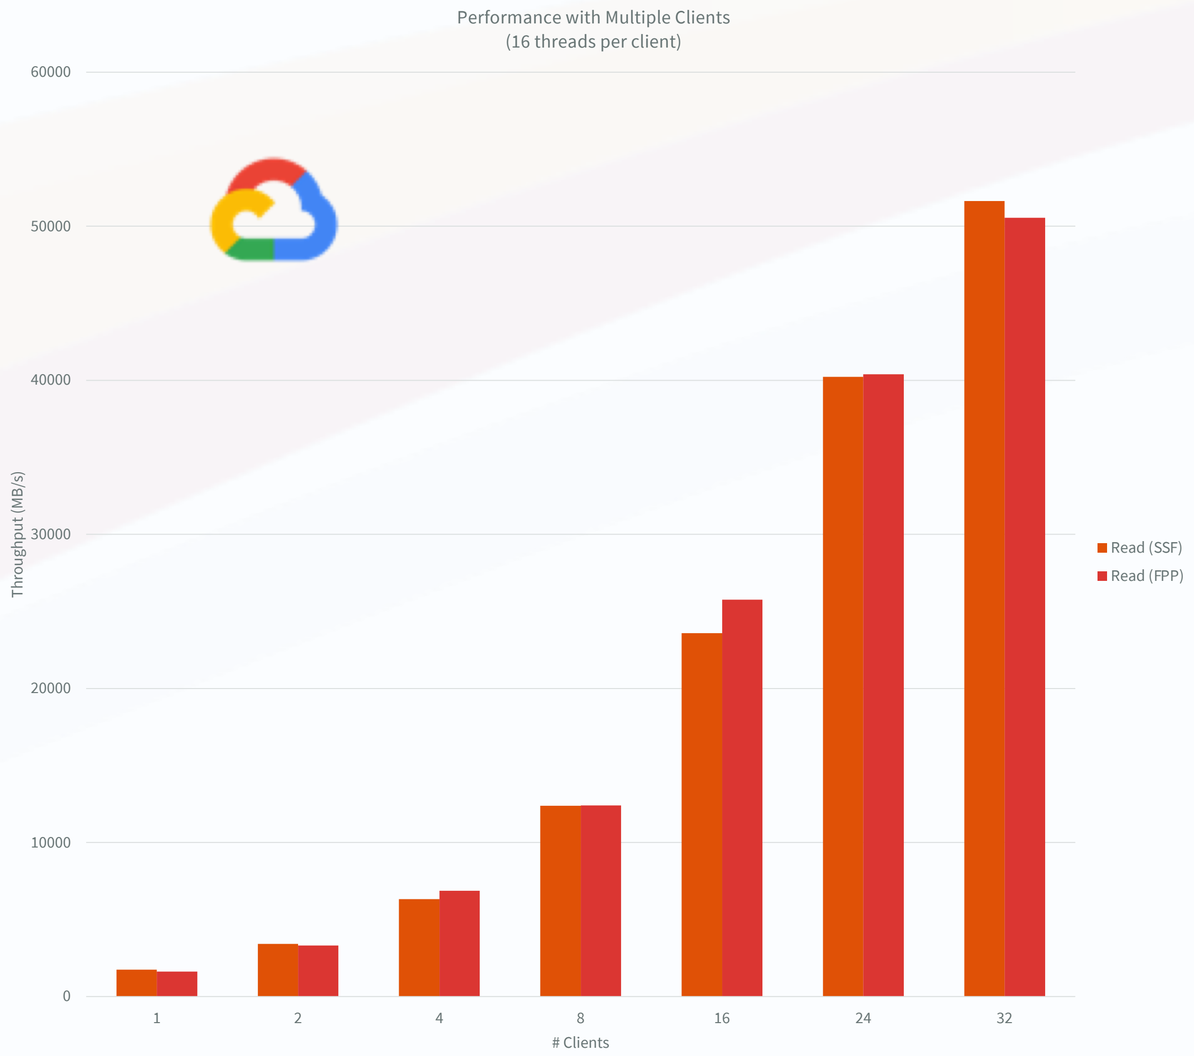

High Performance Kubernetes Support

A³I Performance Direct to Containers

- Kubernetes (K8s) automates deployment, scaling, and management of containerized applications.

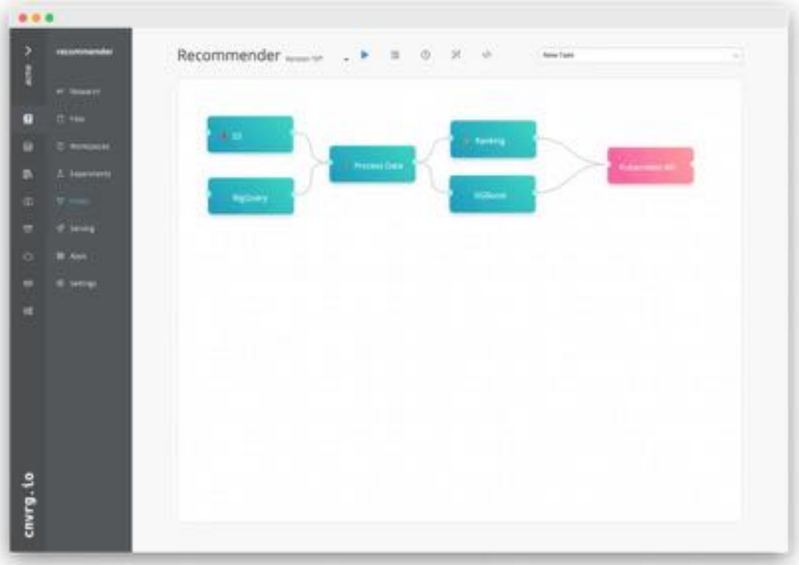

cnvrg.io uses K8s to easily to orchestrate AI workflows. EXAScaler would be seen as a tile

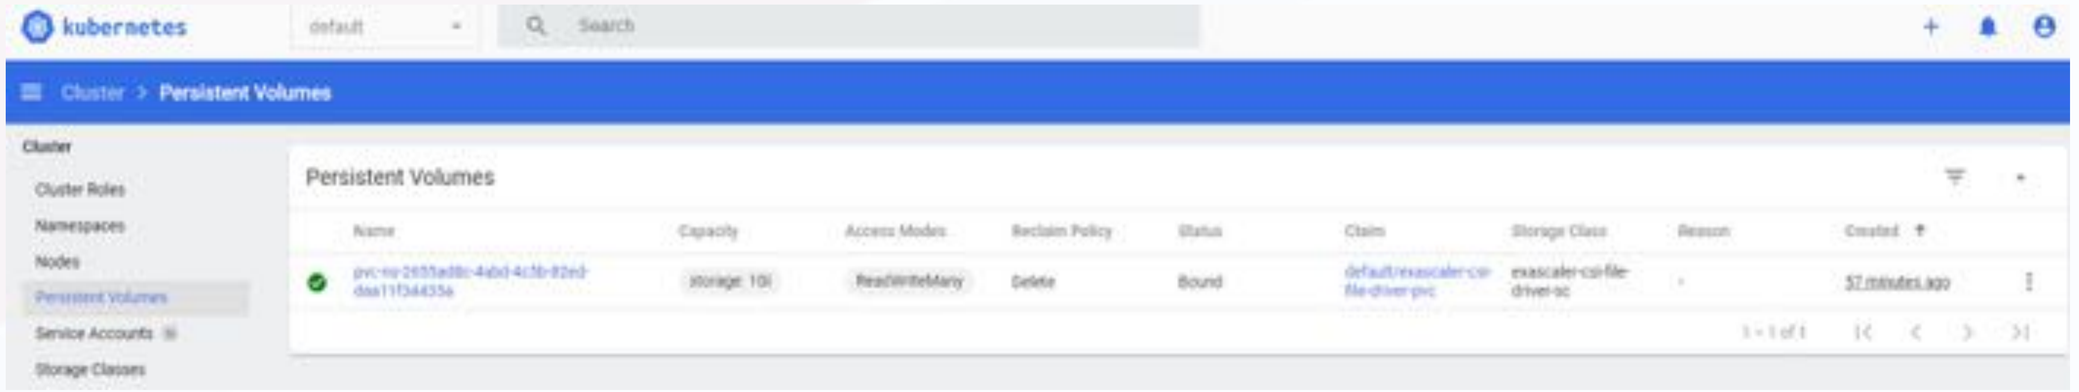

DDN Insight Aggregator for Customers with Many EXAScaler Filesystems

Choose which filesystem(s) performance to view Throughput or IOPs Read, Write or Both

Predict Future Capacity Needs

Breakdown of throughput/IOPs for each filesystem

Capacity Consumed Status for each Filesystem

Dashboard View

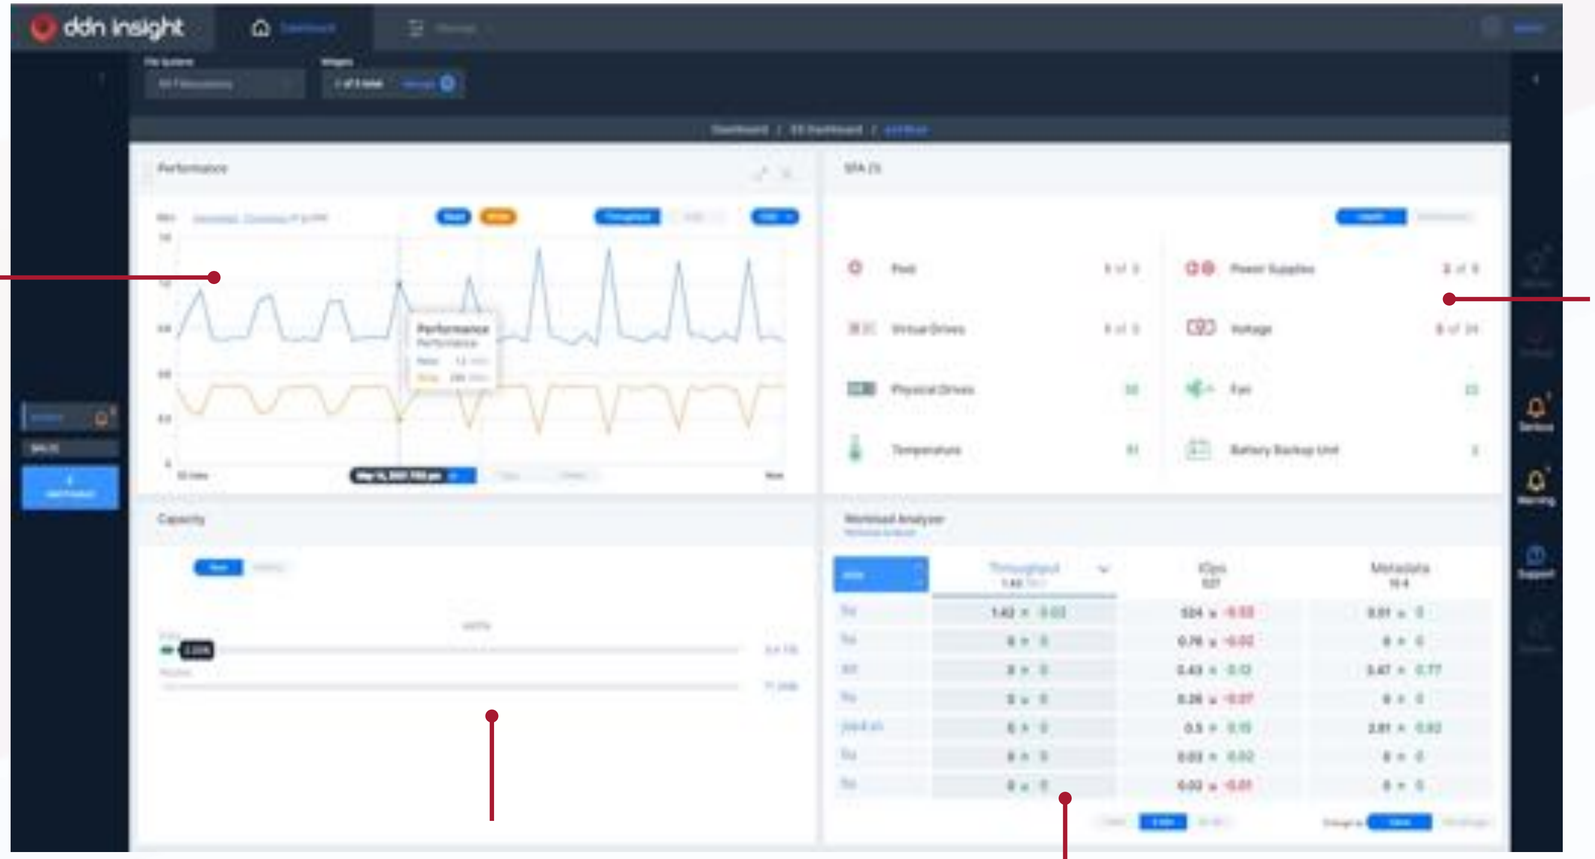

Performance Overview

SFA Health Summary

Capacity Available and Used (Data and Inodes)

Workload Analyzer Widget

EXAScaler Performance

View throughput or IOPs

View Aggregate or Average

Select and De-Select Storage Servers

Read and Write Time Series view

Adjust time view

EXAScaler Performance

Absolute Value and Change Shown

View Metadata Activity

Select all, or a subset of MDS

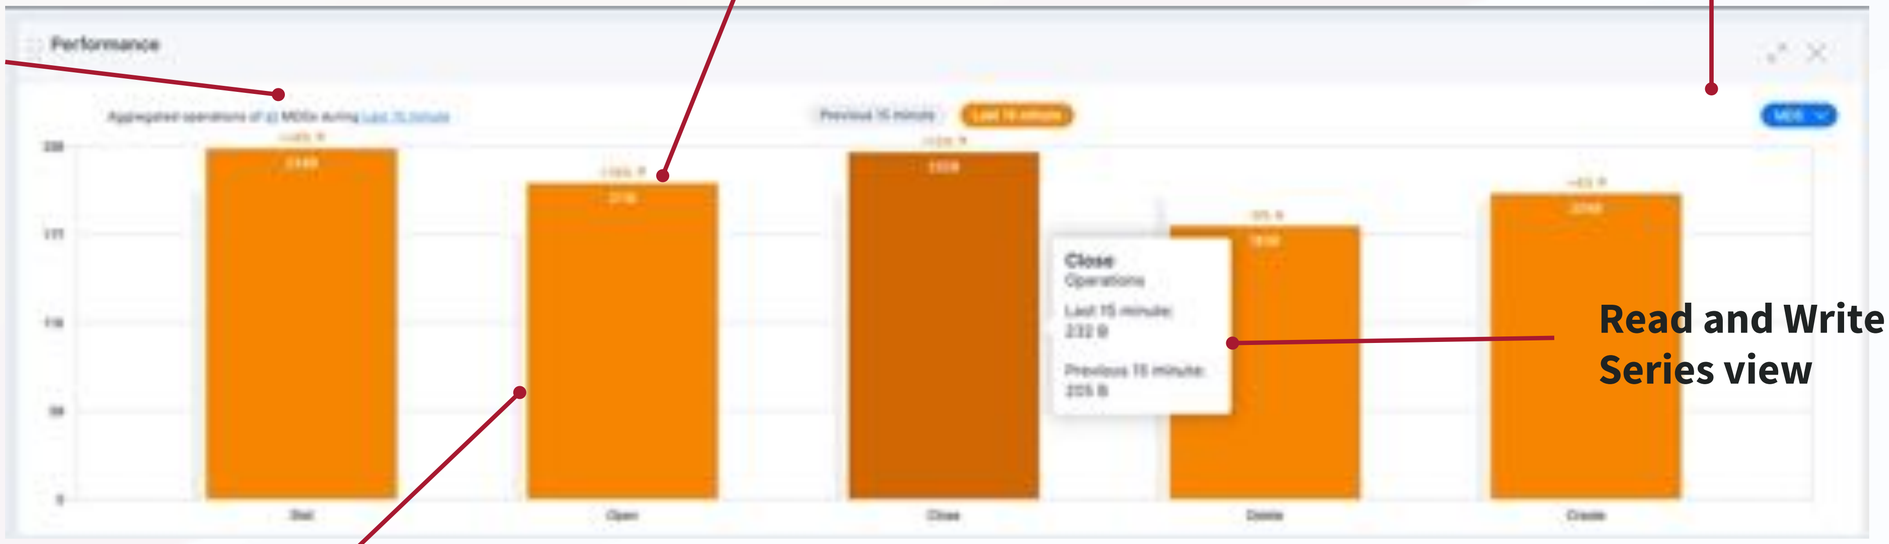

Read and Write Time Series view

Shadow shows performance for the prior time period

Workload Analyzer

User can re-order by Throughput, IOps or Metadata activity

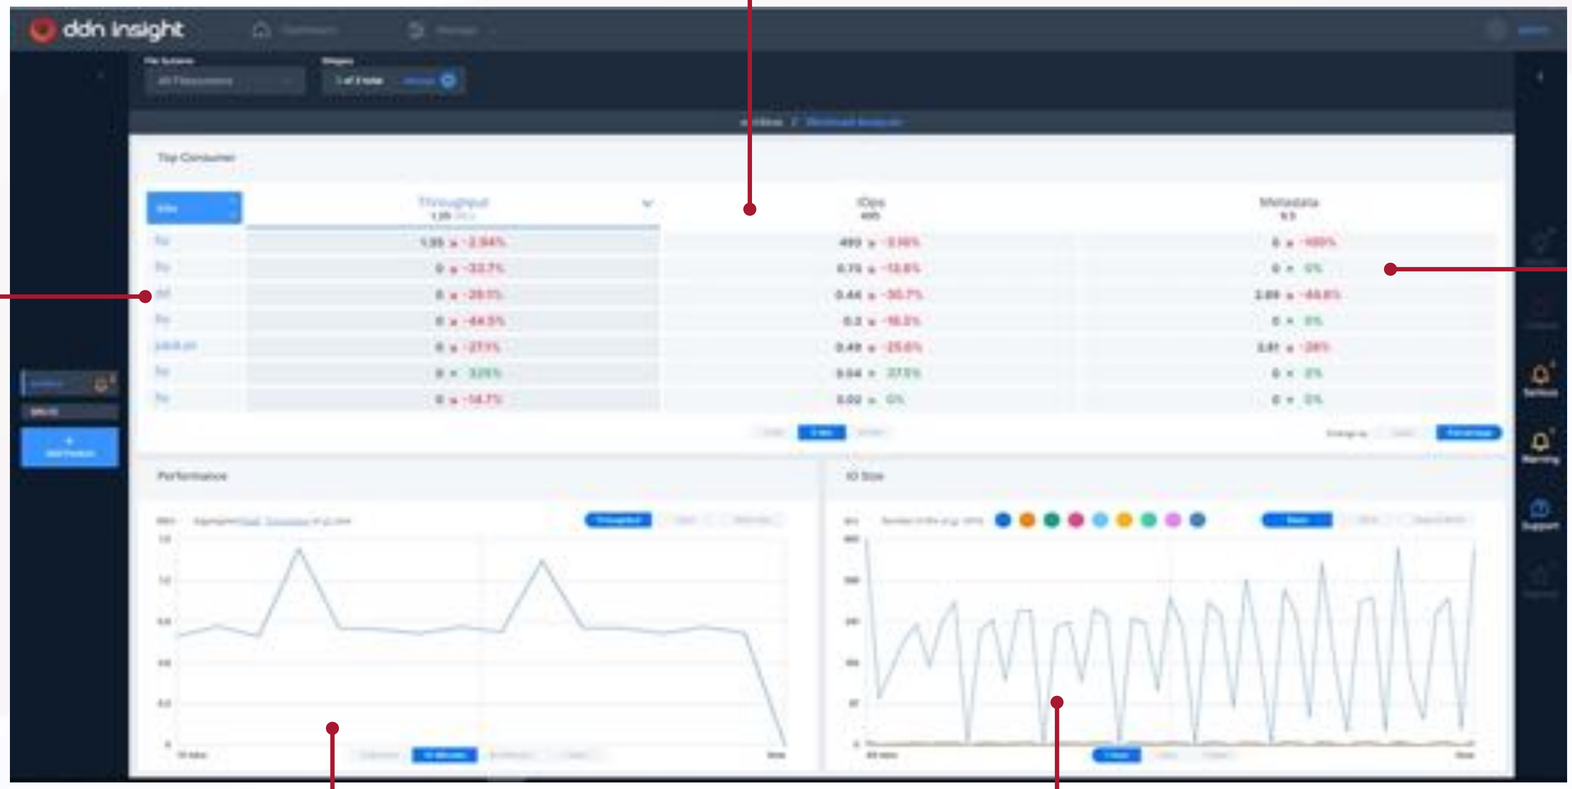

Running Jobs

Workload Analyzer

Aggregate Performance of Selected Jobs

IO Size (read, Write, or Both)

EXAScaler Data Services

- Scale-Out Data Services deliver multiprotocol access to EXAScaler Namespace
- Management REST API for system admin

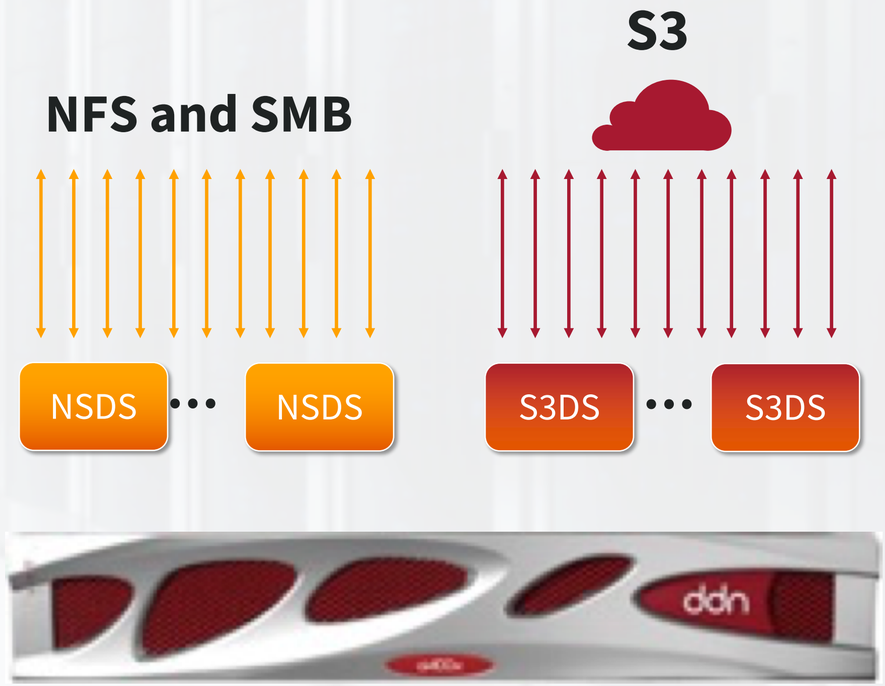

NVIDIA DGX A100 Internal Architecture

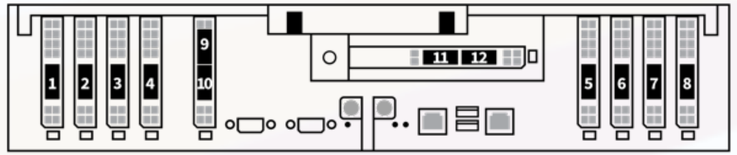

Two single port NICs (cluster) and two GPUs per PCIe switch, up to 24 GB/s per NIC, optimal path for GPUDirect Storage.

Dual port NICs (storage) close to CPU, up to 28 GB/s per NIC.

Eight GPUs interconnected within the system through NVSwitches (NVIDIA proprietary interconnect, 10x higher bandwidth than PCIe Gen4, 600 GB/s GPU-to-GPU).

GPUs from multiple systems can communicate through a cluster network, ideally via single port NICs on same PCIe Switch as GPU (ports 1-8).

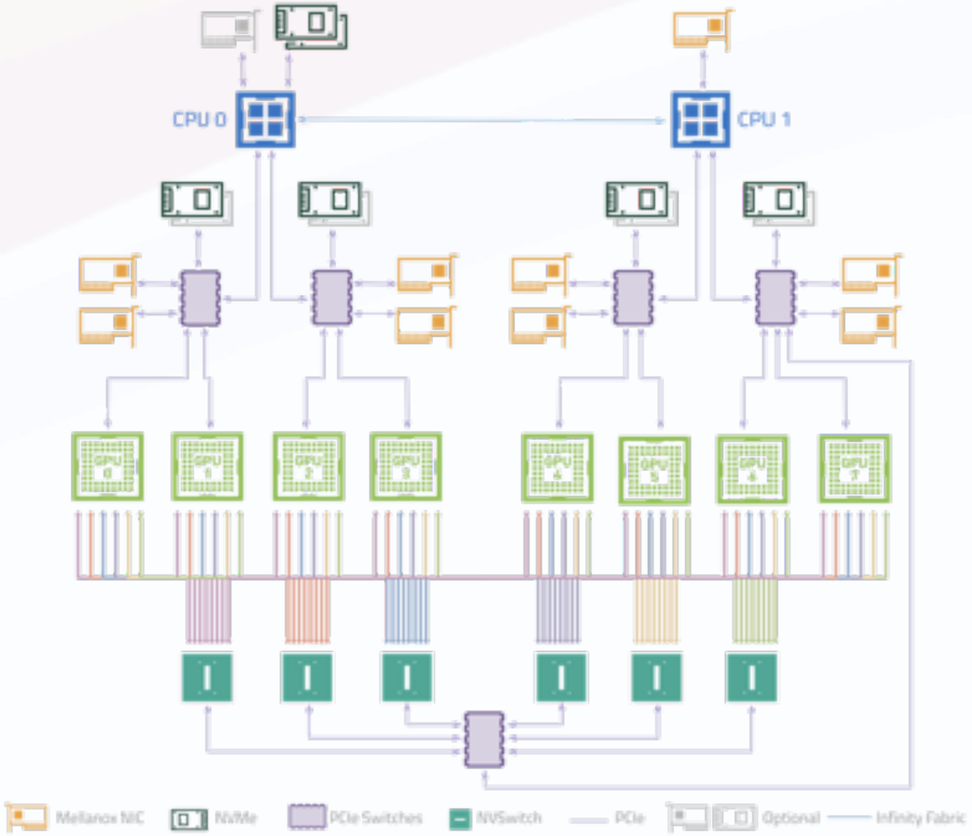

DDN A³I MULTIRAIL ENABLES PLUG-AND-PLAY HPC NETWORKING

FAST, SECURE, RESILIENT NETWORKING MADE EASY

- Multi-Rail automatically detects and manages multiple network interfaces on a single network.
- Load-balancing across network links
- Link fault detection, dynamic failover and recovery
- Improved peak performance with single mount point on client nodes with multiple network interfaces
- Runs over Infiniband and/or Ethernet

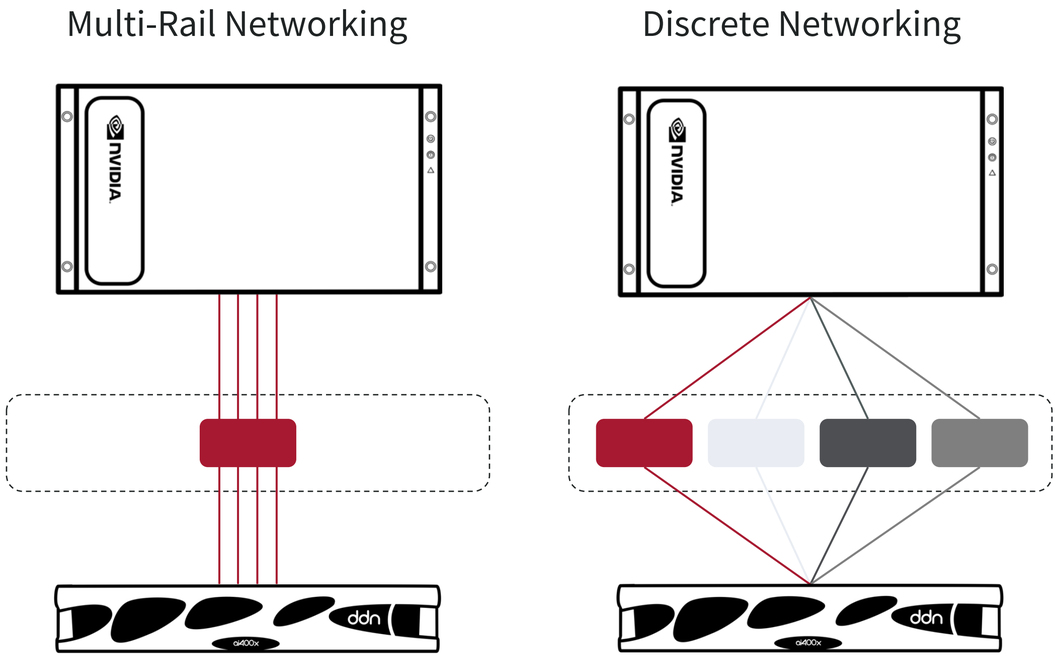

DDN A³I MultiRail greatly simplifies DGX deployments.

UNIVERSAL ACCELERATOR FOR CONTAINERS

DDN A³I enables seamless fastest file access to shared storage directly from containerized applications at runtime, with full performance.

Parallel client capability is inserted at runtime with a universal wrapper and does not require any modification to application or container.

Easy to deploy and manage compartmentalized data access and secure multi-tenancy with trusted levels of segregation using containerized applications.

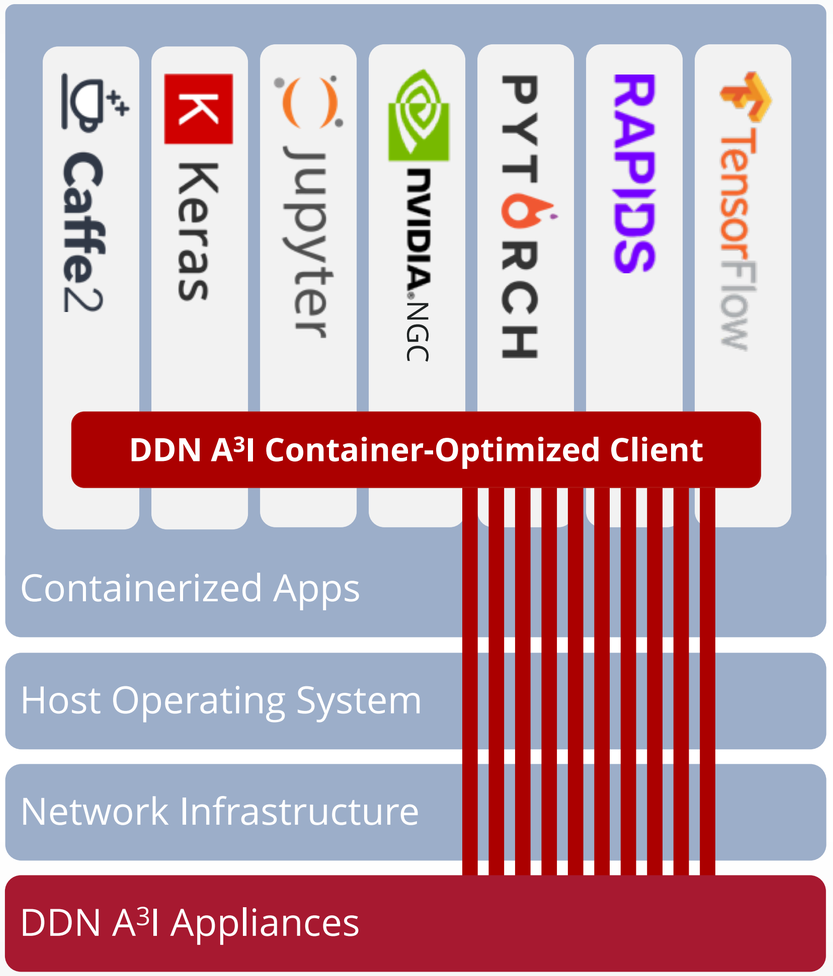

DDN AI400X with NVIDIA DGX A100
Validated Performance Results Summary – Dual Ported NICs

	Two Dual-Port NICs, IB
AI400X Configuration	Single AI400X
DGX A100 Configuration	2 x HDR200 IB Ports 9 and 11
GDS Enabled?	No
Max Read Performance	47 GB/s – 2.7M IOPS
Max Write Performance	41 GB/s
Comments	Optimal for non-GDS deployments

DDN AI400X with NVIDIA DGX A100
Validated Performance Results Summary – Single Ported NICs

	Eight Single Port NIC, IB	
AI400X Configuration	4x AI400X	
DGX A100 Configuration	8 x HDR200 IB Ports 1, 2, 3, 4, 5, 6 ,7, 8	
GDS Enabled?	No	Yes
Max Read Performance	108 GB/s – 4.8M IOPS	178 GB/s
Max Write Performance	95 GB/s	154 GB/s

Accelerating The AI–Enabled Drug Discovery Revolution

Recursion Boldly Reimagines Pharmaceutical R&D With AI

- Leveraging AI to rapidly identify compounds for new treatments.
- DDN and NVIDIA solution based on the DGX SuperPOD architecture, fully integrated and optimized for the Recursion pipeline.
- End-to-end workflow managed; data capture, AI data preparation, all phases of deep learning workloads, and complete data governance, security and migration from the cloud.
- Up to 20X cost savings and introduces new levels of pipeline acceleration.

Empowering GPU-accelerated Clinical and Research Genomics

St-Jude Children's Research Hospital

Leading the Way for The World to Understand, Treat and Defeat Childhood Cancer With AI

- DDN selected to design, implement and optimize the data infrastructure and end-to-end workflow integration.
- Successful and rapid implementation using pre-defined DDN A^{3}I solutions reference architecture enabled quick deployment to researchers with full performance and immediate acceleration for science efforts.

" DDN A^{3}I solutions give us full performance out-of-the-box.

With other storage vendors, our applications run 20X slower and GPUs waste 90% of cycles waiting for data to be delivered. "

Wei Guo
Computational Engineer
St Jude Children's Research Hospital

Accelerating Swedish AI Research across Industry and Academia

Wallenberg Artificial Intelligence, Autonomous Systems and Software Program

- Artificial intelligence and autonomous systems acting in collaboration with humans, adapting to and learning from their environment through sensors, information and knowledge, forming intelligent systems-of-systems.

- DDN A^{3}I All Flash Solutions supports the supercomputer, a NVIDIA DGX SuperPOD, which will deliver a processing speed of 300 petaFLOPS for AI.

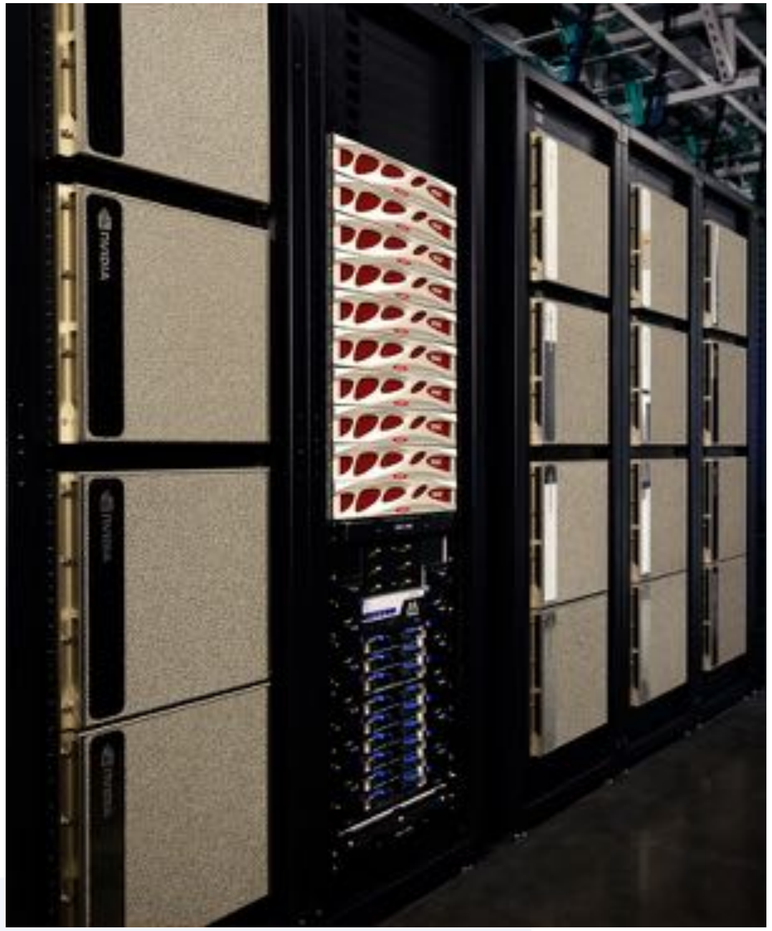

The World's Fastest AI Supercomputer in Academia.

Transforming our Society and Tools with AI

- From religion to agriculture, and liberal arts to engineering, every student will have worked with the AI curriculum and will be able to apply that knowledge to advance the field they are studying in. With performance of an astounding .70 petaflops
- Tightly integrated and tested **A³I and SuperPOD architecture** will strengthen UF research potential through unmatched access to AI training and tools, ultimately addressing some of the most challenging global plights, like rising sea ocean levels, food insecurity, aging populations and more.

"We have a history of success with DDN storage systems powering our HPC computing and anticipate similar high productivity for our AI workloads."

AI Case Study: NVIDIA SELENE

DDN Delivered:

- A dramatic increase in AI workload performance demonstrates how organizations can iterate faster and boost data science productivity
- A modular platform that ensures AI infrastructure scalability with greater speed and cost efficiency
- From delivery dock to fully installed in 4 hours
- A long-term technical collaboration for AI excellence

“Having a partner who stands shoulder-to-shoulder with our engineers to solve the big challenges is where the true value comes from. We’re definitely pushing the boundaries of what’s possible today while exploring new frontiers for the future.”

Throughput on NVIDIA Selene

- Proven Scaling to Reach Max Performance of DDN AI400X systems to 128 NVIDIA DGX100
- MultiRail across both DGX100 systems and DDN AI400X systems

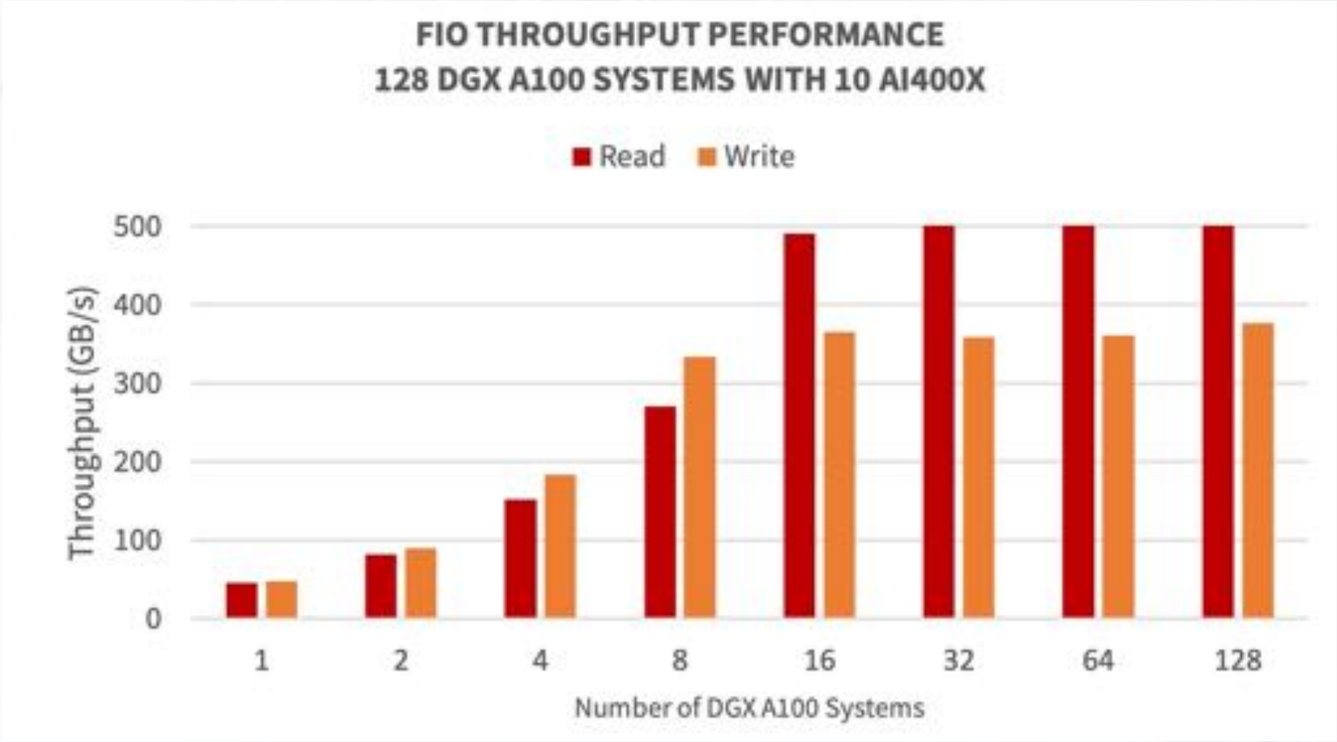

Increased Efficiency, Higher Acceleration, Lower Cost at Any Scale

Proven with thousands of successful deployments

- Wide set of Functionality to bring the right feature set and capabilities to new AI customers, bring easier deployment and monitoring
- Proven at larger scale than any other storage system with the only SuperPOD reference architecture based on EXAScaler
- More to come around simplification and features for our growing customer base.