



Accelerated Data for Financial Analytics

ddn



Powering the Most Demanding Data Intensive Organizations

DDN provides insight into data and innovation around storage to the world's largest enterprises and government agencies.

HPC-like workflows in the Enterprise market, driven by AI and huge analytics, are moving IT organizations to embrace parallel file systems and scalable data management platforms.






DDN Storage Ranks Number One in Intersect360 Storage Survey Ahead of All Systems Companies, Suppliers and Startups

“DDN has a strong technical perception and broad usage, with good markers for both loyalty and future adoption.”

Storage: Technical Impression





NVIDIA SELENE SUPERPOD

Powered by DDN AI400X



DDN A³I solutions are proven at-scale on the largest SuperPOD with DGX A100 in operation globally:

- 40 DDN AI400X appliances
 - ~2TB/s throughput, 120M IOPS
 - Fully-integrated and optimized for DGX A100 systems
 - Rapid deployment, predictable phased scaling
- 280 NVIDIA DGX A100 systems
- HDR 200 InfiniBand Storage Network

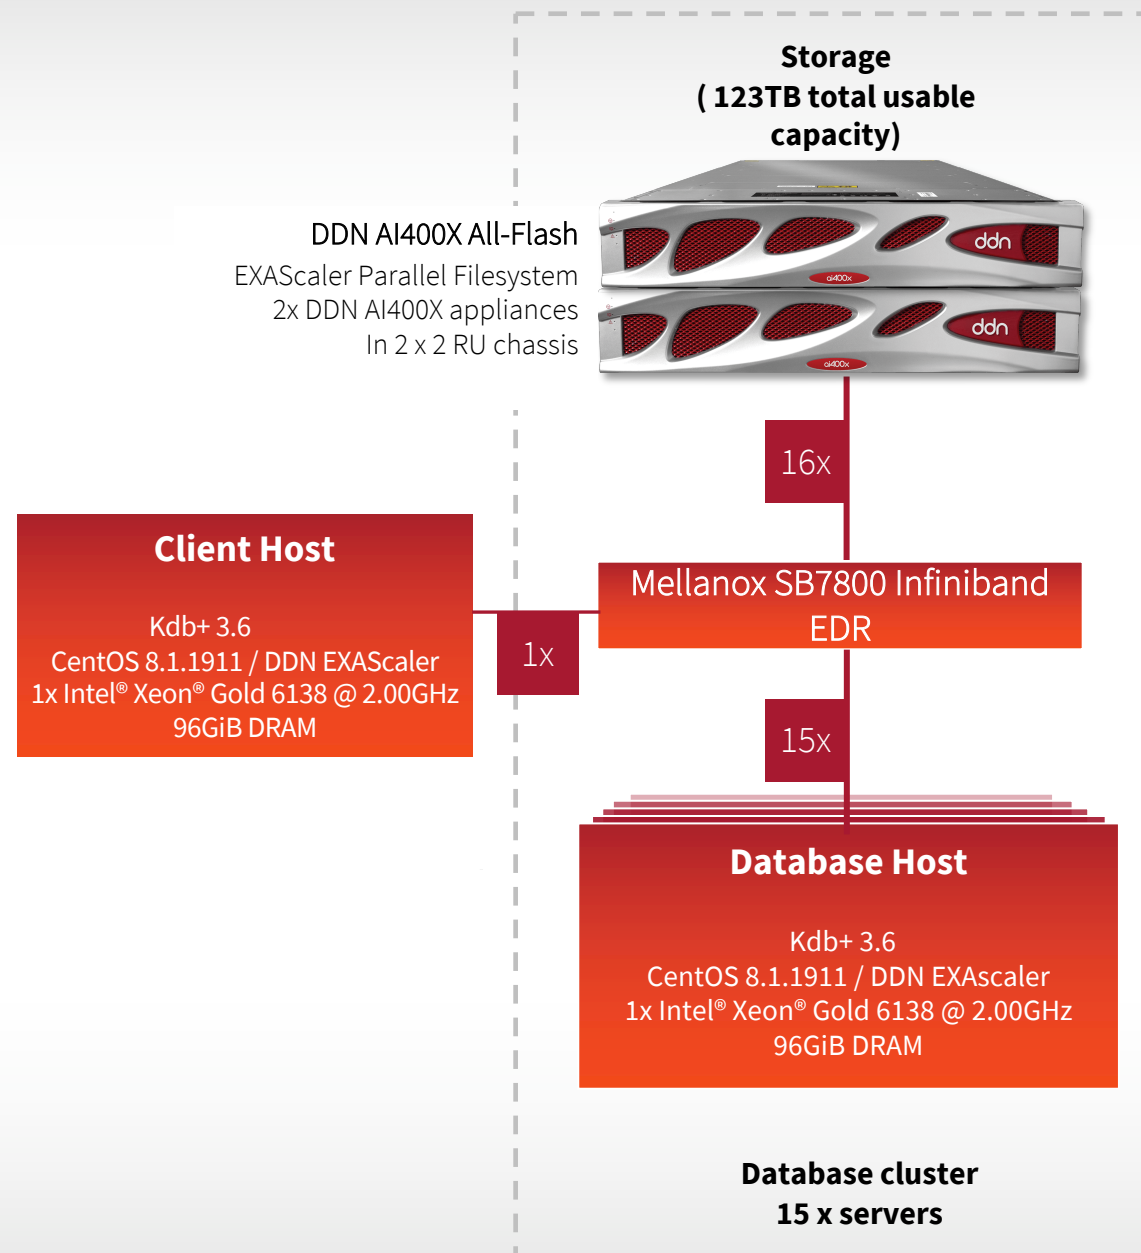
Fully-optimized, easy to deploy solutions with AI400X available for DGX A100, POD and SuperPOD at all scale.





DDN AI400X Test Platform

- Simple All NVMe DDN Storage Appliance with a complete parallel filesystem embedded
- KDB+3.6 Time Series Analytical Database Software
- Total 16 client systems (database Hosts and one client hosts)





DDN with 15 server nodes vs Other Storage Solution with 14 DB server nodes

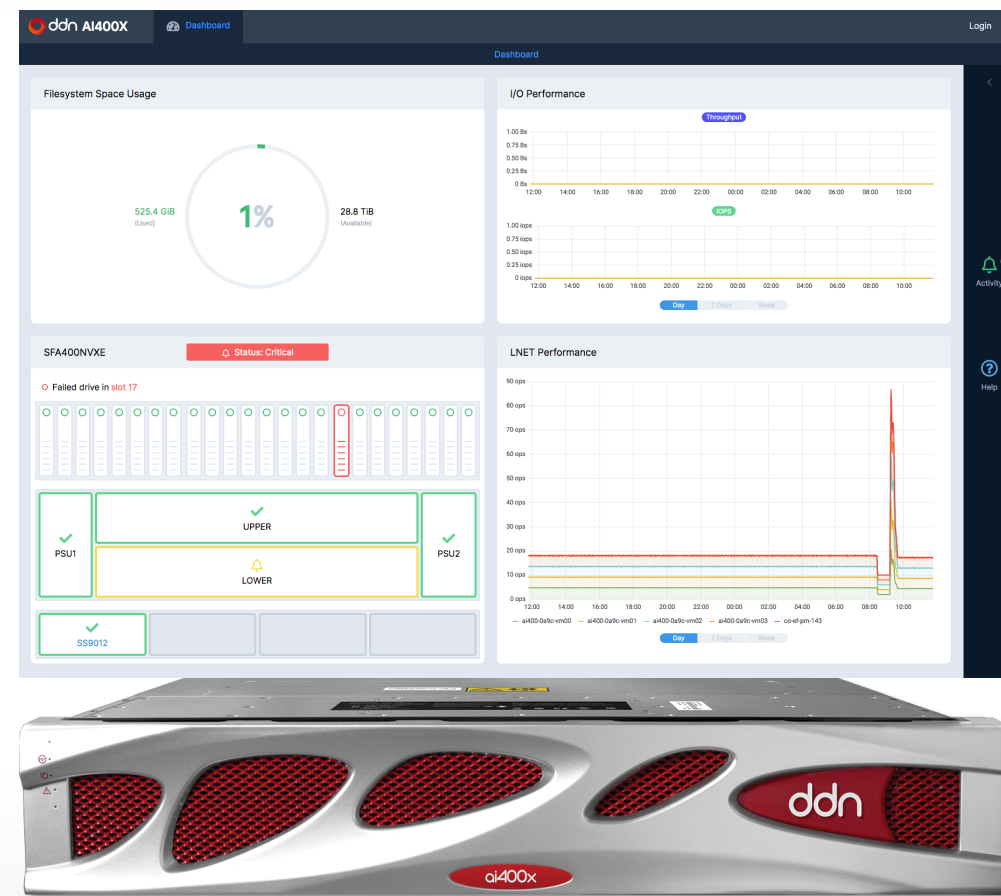
Versus a kdb+ solution based on a parallel file system with 14 database server nodes and 18 storage server nodes (SUT ID KDB200401):

- *was faster in all four **10-user market snapshot** mean response time benchmarks in the Kanaga suite, with speedups ranging from 1.8x for year 2 to 2.0x for year 4 (STAC-M3.β1.10T.YR[n]-MKTSNAP.TIME)*
- *was faster in 4 of 17 mean response time Antuco benchmarks, including:*
 - *9.4x speedup in **NBBO** (STAC-M3.β1.1T.NBBO.TIME)*
 - *46% faster in **10-user market snapshot** (STAC-M3.β1.10T.MKTSNAP.TIME)*
 - *15% faster in **10-user aggregate stats** (STAC-M3.β1.10T.STATS-AGG.TIME)*
 - *3% faster in **100-user interval stats** (STAC-M3.β1.100T.STATS-UI.TIME)*
- *was within 0.2% of the speed in the 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)*



The Ideal Choice for Fast Financial Analysis






- Mature and High Performance Appliances that start at 2RU and Scale up or Scale Out
- Easy Setup, speedy operation
- Performance for a single 2RU up to 48GB/s for reads and 3M IOPs






Performance, Simplicity and Flexibility

Configuration	AI400X	+ 1 SS9012	+2 SS9012	+4 SS9012
NVMe	24	24	24	24
3.5" SAS	0	90	180	360




ACCELERATED AI AND HPC SOLUTIONS FULLY-OPTIMIZED FOR NVIDIA DGX



MAX PERFORMANCE




COMPLETE WORKFLOWS




LIMITLESS SCALING




DATA MANAGEMENT



MULTI-TENANCY



DATA SERVICES



GPUDIRECT TO STORAGE*



REAL-TIME ANALYTICS



WHY DO CUSTOMERS **PREFER DDN?**



FULL PERFORMANCE PROVEN AT SCALE

Our shared parallel architecture ensures full GPU saturation on even the largest DGX SuperPOD with DGX A100 currently in operation.



RELENTLESS FOCUS ON AI AND HPC

Our leadership comes from long-term investment in leading-edge research and development that solves our customers' end-to-end data needs.



ABILITY TO EXECUTE TRUSTED GLOBALLY

We have over 20 years of experience and expertise delivering solutions to the most demanding organizations worldwide, including NVIDIA.



ddn