



DDN Partner Update

Spectrum Scale Usergroup Meeting SC18

November, 2018

Sven Oehme – Chief Research Officer DDN
oehmes@ddn.com

DDN GRIDScaler versions

	Availability				
	GS200NV	GS400NV	GS7990	GS14KX	GS18K
GRIDScaler v4	✓	✓*	✓	✓	✓*
v4 upgrade to v5	✓	✓*	✓	✓	✓*
GRIDScaler v5	✓	✓*	✓	✓	✓*

* Upon platform availability



Optimizations for GRIDScaler V5

Optimizations for GRIDScaler V5

- ▶ Updated device drivers, OS and Scale tuning parameters and SFA multi-queue LUN support
- ▶ Embedded systems can now achieve up to 1.2 Million random 4k read IOPS
- ▶ External SFA14KX NSD Server performance went from 1.25 Million to 2.96 Million*
- ▶ These enhancements were used to produce the SpecSFS 2014 record publications


* test was using external NSD Server. all numbers are measured from network attached clients with GPFS PERF using one 100 GB file per client during random 4k reads using O_DIRECT

DDN SFA14KX

Fastest, Densest and Simplest at Scale

Low Latency, Highly Efficient Architecture

- All in one integrated design with expansion capability
- Dual Redundant Controllers
- 72 Drive High-Density 2.5" Enclosure with NVMe support for 48 2.5" dual ported NVMe
- Optimized Building Block for BW or IOPs
- Support for up to 20 SS9012 12Gb/s 90 drive Enclosures



Flexible Connectivity

- ▶ 10/40/100 GbE
- ▶ OmniPath
- ▶ 16/32Gb FC

Industry Leading Performance

- ▶ 72 SAS SSD or 48 NVMe
- ▶ Up to 60 GB/sec throughput
- ▶ Up to 4 million IOPS

Best Data Protection


- ▶ Fully Declustered RAID
- ▶ Higher Data Availability
- ▶ Flexible Pool Management
- ▶ Optimized for both Random and Sequential IO



World record SpecSFS 2014 with GRIDScaler


DDN SFA14KX with GRIDScaler

- ▶ With the SFA14KX and GRIDScaler parallel filesystem, DDN gains pole position for SPEC SFS
- ▶ DDN's SFA14KX running SFAOS with Declustered RAID and connecting to 6 GRIDScaler servers Sustains 25% more builds at 1/3rd of the Overall roundtrip latency with read-intensive SSD compared to previous #1 with Optane NVMeof - the next nearest competitor




System Benchmarked for SPEC SFS

SpecSFS 2014 – SWBUILD compare to other vendors




25% more builds at 1/3rd of the latency with read-intensive SSD compared to previous #1 with Optane NVMe

SpecSFS 2014 – VDI compare to other vendors




SpecSFS 2014 – DATABASE compare to other vendors

SpecSFS 2014 - DATABASE



SpecSFS 2014 – EDA compare to other vendors


SpecSFS 2014 - EDA





DDN A³I Solutions: Turnkey, integrated and optimized
for NVIDIA DGX-1 and HP Apollo 6500

SCALE UP, SCALE OUT OR SCALE BOTH




DDN A³I SOLUTIONS TO A SINGLE CONTAINER ON DGX-1

23 GB/s and 395K IOPS to a single container*

DDN A³I parallel storage client demonstrates over 23 GB per second and over 395K IOPS to a single container on DGX-1.

Typical deep learning codes perform IO using 128K size for which DDN delivers over 20 GB/s of sustained performance.



*numbers are with a single AI200 and was limited by client side performance of single DGX client


DDN A³I SOLUTIONS TENSORFLOW TRAINING PERFORMANCE

Fast, Consistent, Linear AI and DL Performance

DL Training application performance scales linearly using multiple GPUs on DGX-1 with DDN parallel storage.


Parallel storage performance and shared architecture magnify end-to-end DL workflow acceleration.

Extensive application interoperability and performance testing has been engaged by DDN in close collaboration with NVIDIA and customers.




DDN A³I SOLUTIONS LEADS PERFORMANCE FOR AI AN DL

ResNet-152



ResNet-50



The Register
Biting the hand that feeds IT

"In the Resnet-152 and Resnet-50 tests, the AI200 tested faster than competing Pure, NetApp and Dell EMC systems."

Systems being installed now



SpecSFS 2014 for AI ??

DDN started an effort with multiple other vendors to work on a standardized i/o test framework for AI/ML workloads we would like to be releases by the SPEC organization. If you want to help with some strace snippets from your workload please talk to me after the presentation or send me an email oeames@ddn.com !

Come see us at our Booth!

► Demos on

- A3I
- IME on NFS
- IME I/O Through Node ejection
- Data Management
- Super Secret things





[DDN.COM/A3I](https://www.ddn.com/A3I)

Thank You!

Keep in touch with us.



sales@ddn.com



9351 Deering Avenue
Chatsworth, CA 91311



[@ddn_limitless](https://twitter.com/ddn_limitless)



1.800.837.2298
1.818.700.4000



[company/datadirect-networks](https://www.linkedin.com/company/datadirect-networks)