


DDN 2018

DataDirect Networks Japan
Hideaki Fukui



DDN | A Global Company For All Your Data Storage & 20 years

1998 DDN US Start

2008 DDN Japan Start

**2018 DDN 20th anniversary of the founding
DDN Japan 10th anniversary**

World's Largest Private
Global Storage Company

Delivering Exabytes of
Valuable Data to
Thousands of the
World's Largest
Organizations






Solving Complex On-
Prem and Cloud Data-At-
Scale Challenges

DDN's Global Field, Support and
Engineering Teams Cater to Your Data
Challenges and at Any Scale



DDN | Scalable Storage Solutions for Your Specific Applications

- ▶ **Performance** Storage
Platforms designed for scale
- ▶ All Aspects of **Scale**: IOPs, Throughput, Latency, Capacity, Distribution
- ▶ **Flexible** Architecture Design
- ▶ Global Services and Support from **Storage Specialists**



World-Renowned & Award Winning

DDN | DDN Product Offering



Storage Software

- WOS
- IME
- EXAScaler SW



SDS Appliances

- WOS
- IME
- EXAScaler ZFS



SFA Platforms

- SFA Block
- SFA Virtualization
- ES Appliances
- GS Appliances




Other Platforms

- JBODs
- Monitoring
- Special Development Projects

DDN | Maximize the Value of Data at Scale

Top Performance

We have delivered in production the fastest and most efficient NVMe and data storage platforms – over 1TB/s and 50M IOPs.



True Flexibility

NoSQL database? Fraud Detection? Analytics at scale? DDN has successfully met the needs of our customers' most challenging workflows.



Seamless Scale

DDN's solutions provide linear scaling from midrange on premise needs to your largest scale globally distributed file system and object requirements.



Customer Satisfaction


DDN's worldwide team of technical experts has only one mission – delighting all our customers in all the markets we serve.





Product Update

Flash cost performance “Down” per GB [2013 - 2018]



DDN | Full Refresh of the DDN Product Line in 2018

- ▶ Flash Solution SFA200NV
- ▶ Midrange SFA7990
- ▶ High Density JBOD SS9012


- ▶ Twice as Fast, Same Cost
- ▶ New Flash Platforms
- ▶ Many Additional Features
- ▶ Margin Accretive

Flash & NVMe Platforms

SFA200NV



SFA200NVE




Hybrid Data Platforms, Integrated File System Appliances

SFA7990



SFA7990E



SS9012





**FLEXIBLE ALL
NVM_e FLASH**

20GB/s



Parallel, Flexible Flash
Performance

DDN SFA200NV

24x NVM_e

800GB- 1.6/3.2/6.4 TB

960GB- 1.92/3.84/7.68/15.36 TB

I/F



4x IB EDR VPI

2x OPA (Embedded Parallel Fileserver)

8x 8GbpsFC

8x 16/32GbpsFC

AI200 (RANDOM READ IOPS AND THROUGHPUT)





**FAST & SCALABLE
HYBRID STORAGE
INTEGRATED FILESYSTEMS**

20GB/s



Parallel File System
Performance & Capacity with
Hybrid Flash + Disk Drives

DDN SFA7990

4U 90x SSD / NL-SAS HDD

800GB- 1.6/3.2/6.4 TB

960GB- 1.92/3.84/7.68 TB

4.0 / 6.0 / 8.0 / 10.0 / 12.0 / 14.0 TB

I/F


4x IB EDR VPI

2x OPA (Embedded Parallel Fileserver)

8x 16GbpsFC

DDN SS9012


- ▶ SS9012 will be the DDN standard drive enclosure
- ▶ Supported on SFA14KX, and SFA7990 series products



SS9012	
RU	4
Disk Slots	90
PSUs	4 PSUs (2+2 redundant)
Cooling	5 independent fans
IO Modules	x2 IO modules, Fully 12Gb/s SAS x4 4-Lane Ports per IO Module
Baseboard	12Gbps throughout
Cabling	Copper or optical SAS

SFA7990 Performance Data

IOR Sequential Throughput of AI7990





SFA7990 at Scale by SS9012 / 14TB HDD

- **Start at 4U and Scale-out or Scale-Up**

Shared data access through a unified, scalable namespace. Add HDD or SSD to a single SFA7990

- **Flexible Implementation and Performance**

An all-round performer across Throughput and IOPs, File and Block, SSD and HDD



180' HDD
18-20 GB/s
2.0 PB

270' HDD
20+ GB/s
3.0 PB

450' HDD
20+ GB/s
5.0 PB

DDN | Into the enterprise AI mainstream by A3I solutions

- ▶ The AI200 has up to 360TB of 24 x dual-ported NVMe flash drive capacity inside its 2U enclosure, and hooks up to the DGX-1 with either 4 x EDR InfiniBand (EDR IB) or 100Gbit/s Ethernet (100 GbitE). It delivers up to 20GB/sec of file system sequential read throughput and over 1 million IOPS.
- ▶ The larger AI7990, in a 4U cabinet, reverts back to 20Gb/sec sequential read performance and provides up to 700,000 IOPS. It supports 90 x 3.5-inch slots for SSDs and disk drives. There can be up to 4 expansion chassis, each with 90 bays, providing up to 5.0PB of usable capacity.



DDN A³I WITH DGX-1 LAUNCH PARTNERS





SFAOS Update


SFA OS11.x New feature



- Declustered RAID
- NVMe Support
- Multiple VD per pool / Dual port Drive access
- Strip Groupe Count (DCR IO perf)

Embedded platform:

- HCA PF Pass-thru
- VirtIO-SCSI driver



Declustered RAID

SFAOS FAST, DECLUSTERED OPERATION



HIGHER AVAILABILITY

Faster Rebuilds, and more options for sparing and spare capacity management




MAINTAINS FLEXIBILITY

Highly Flexible Pool Management allows simple creation of VDs across any-size HDD/SSD Pools



PERFORMANCE OPTIMIZED


IOPs optimized to deliver the most IOPs to the platform – Block IOPs AND File IOPs



Legacy RAID vs DCR

Legacy RAID - 31Drive

3x RAID6(8D+2P) +1 Spare Drive




Example with 3 RAID 6 (8+2) pools

Rotating 8+2 parity across the Pool (or more correctly, across the VD, but here we have 1:1 VD:Pool)

RAID Pools has a RAID stripe size which is defined by the chunk size – write one 8 chunks of data and 2 chunks of parity for each stripe separate hot spares are optionally available in the event of a drive failure

DCR POOL - 31Drive



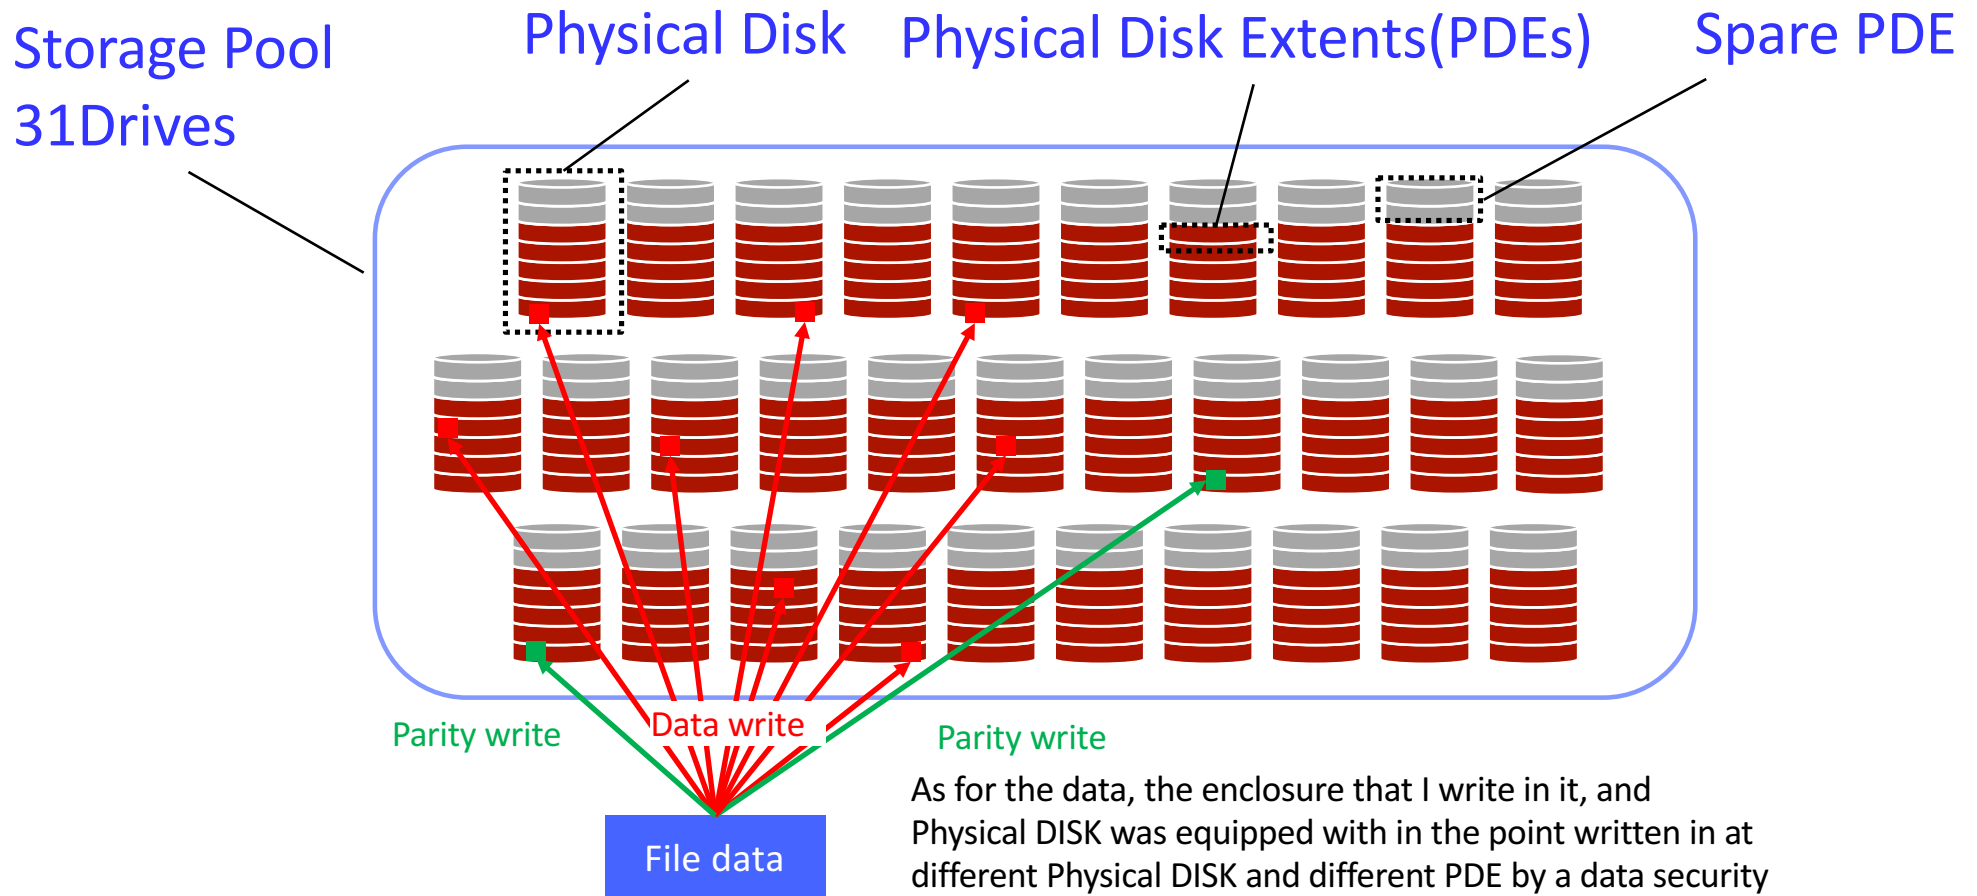
Example with the same number of drives (31)

Now we create a single DCR Pool across all 31 drives

The DCR creation divides all PDs into 8GiB PDEs

The formattable capacity appears lower than traditional RAID as we include hot spare PDEs within the pool (at least large enough to cope with a single drive failure)
e.g. for 31x 8TiB drive we utilise 8TiB for spare PDEs


De-clustered RAID (DCR) Data & Parity Write



As for the data, the enclosure that I write in it, and Physical DISK was equipped with in the point written in at different Physical DISK and different PDE by a data security method of 8D+2P is considered


Rebuild technology

Legacy RAID - 31Drive
3x RAID6(8D+2P) +1 Spare Drive




- Because one drive takes charge of all write processing when rebuild does data using a hot spare drive, it becomes the bottleneck of the rebuild processing performance. Therefore, Rebuild becomes it for a long time, and the risk of the Taro strike increases on the day by the simultaneous obstacle of plural drives. In addition, a performance decreases with becoming it for a long time for long time of rebuild.

DCR POOL - 31Drive



- Many other drives contribute to PDE rebuild
- other PDEs that resided on the failed drive are rebuilt (such that the rebuilding PDEs are spread across multiple drives)

DCR rebuild performance (use 4TB NL-SAS)



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)