DDN Update

UK Spectrum Scale User Group Meeting

Dec, 2019

Sven Oehme – Chief Research Officer DDN

DDN® AT SCALE ENTERPRISE

YOUR STRATEGIC PARTNER FOR MODERN DATA WORKLOADS

ENTERPRISE

Virtualized Enterprise
Software Defined Unified Storage
High Performance Unified Storage

AT SCALE

AI Data Management
HPC & Big Data Storage

At Scale Data Services for
AI, Big Data and HPC

Optimized AI Platforms For
Every Use Case

DDN® AT SCALE ENTERPRISE

High Performance Flash &
Hybrid Unified Storage

SW Defined Storage for Telco
5G, IoT and the Enterprise

Simplicity and Control for
Virtualized Environments

COMPLETE NEW RANGE OF APPLIANCES

- All DDN systems updated with more powerful subsystems and larger cache
- Both improved performance, and increased capability in our Scaler Appliances
- SFA400NVX Delivers over 63M IOPs from NVMe in one Rack
- SFA7990X Delivers 140GB/s from HDD in one Rack
- SFA18KX brings the most powerful Hybrid NVMe-HDD platform supporting up to 1760 Drives

SFA200NVX

SFA400NVX

SFA7990X

SFA18KX

ALL NVME

HYBRID

DDN SFA X-APPLIANCES

SFA200NVX

24r/20w GB/s
1.5M IOPS

UP TO 24 DRIVES

SFA400NVX

48r/30w GB/s
>3M IOPS

UP TO 384 DRIVES
+0 +1 +2 +4

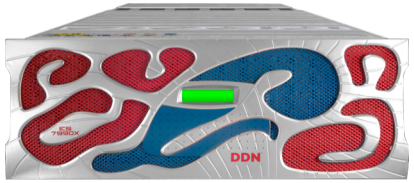

SFA7990X

24r/20w GB/s
800K IOPS

UP TO 450 DRIVES
+0 +1, +2, +4

SFA18KX

70r/70w GB/s
>3M IOPS

UP TO 1872 DRIVES
+0 +5, +6, +8, +10, +16, +20

ALL NVME

HYBRID

SFA400NVX GAINS HDD EXPANSION

- New NVMe Hybrid Platform in a smaller package with the power you need to support complete data services

Configuration	SFA400NVXE	+ 1 SS9012	+2 SS9012	+4 SS9012
NVMe slots	24	24	24	24
3.5” HDD or SAS SSD Slots	0	90	180	360

Expansion Options

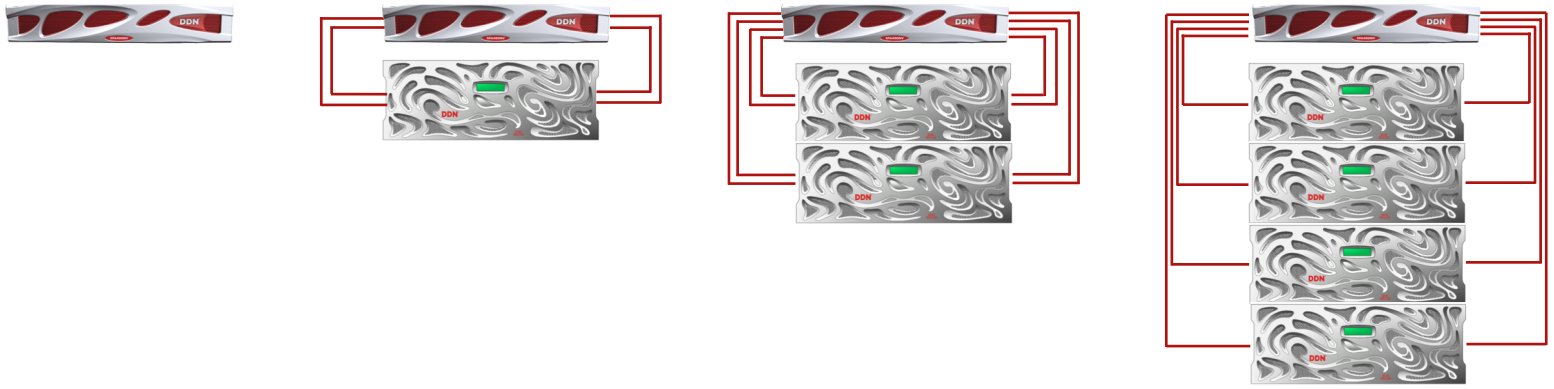

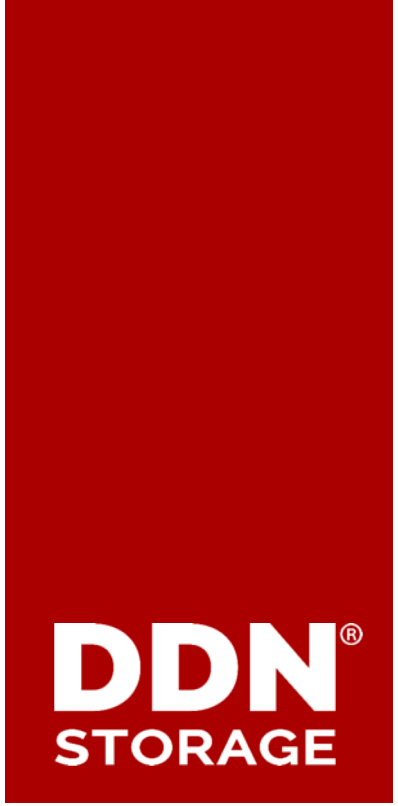

DDN GRIDScaler Platform Optimizations

DDN | GRIDScaler

Massively Scalable NAS & Parallel File Storage Appliance

- Easy to deploy, All-in-One Appliance for All Flash Array with HDD, archive and cloud tiering options
- Scale-out building blocks architecture
  - Configurations scale from <100 TB to PBs of storage and 10s of TBs/sec of performance
- Flash Centric Architecture - custom embedded fabric delivers optimal SSD performance
- Feature-Rich, Enterprise Grade Quality and High Availability with no single point of failure
- Simple, Intuitive but Powerful DDN Insight monitoring solution

LOW LATENCY DESIGNED-IN

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

IO PATHS

TRADITIONAL

Components Simplified with SFA™ Embedded Appliances

SFA EMBEDDED FILESERVICE

CLIENT → HCA/NIC → SWITCH

SCALER™ APPLIANCES

Collapse of layers improves simplicity and performance

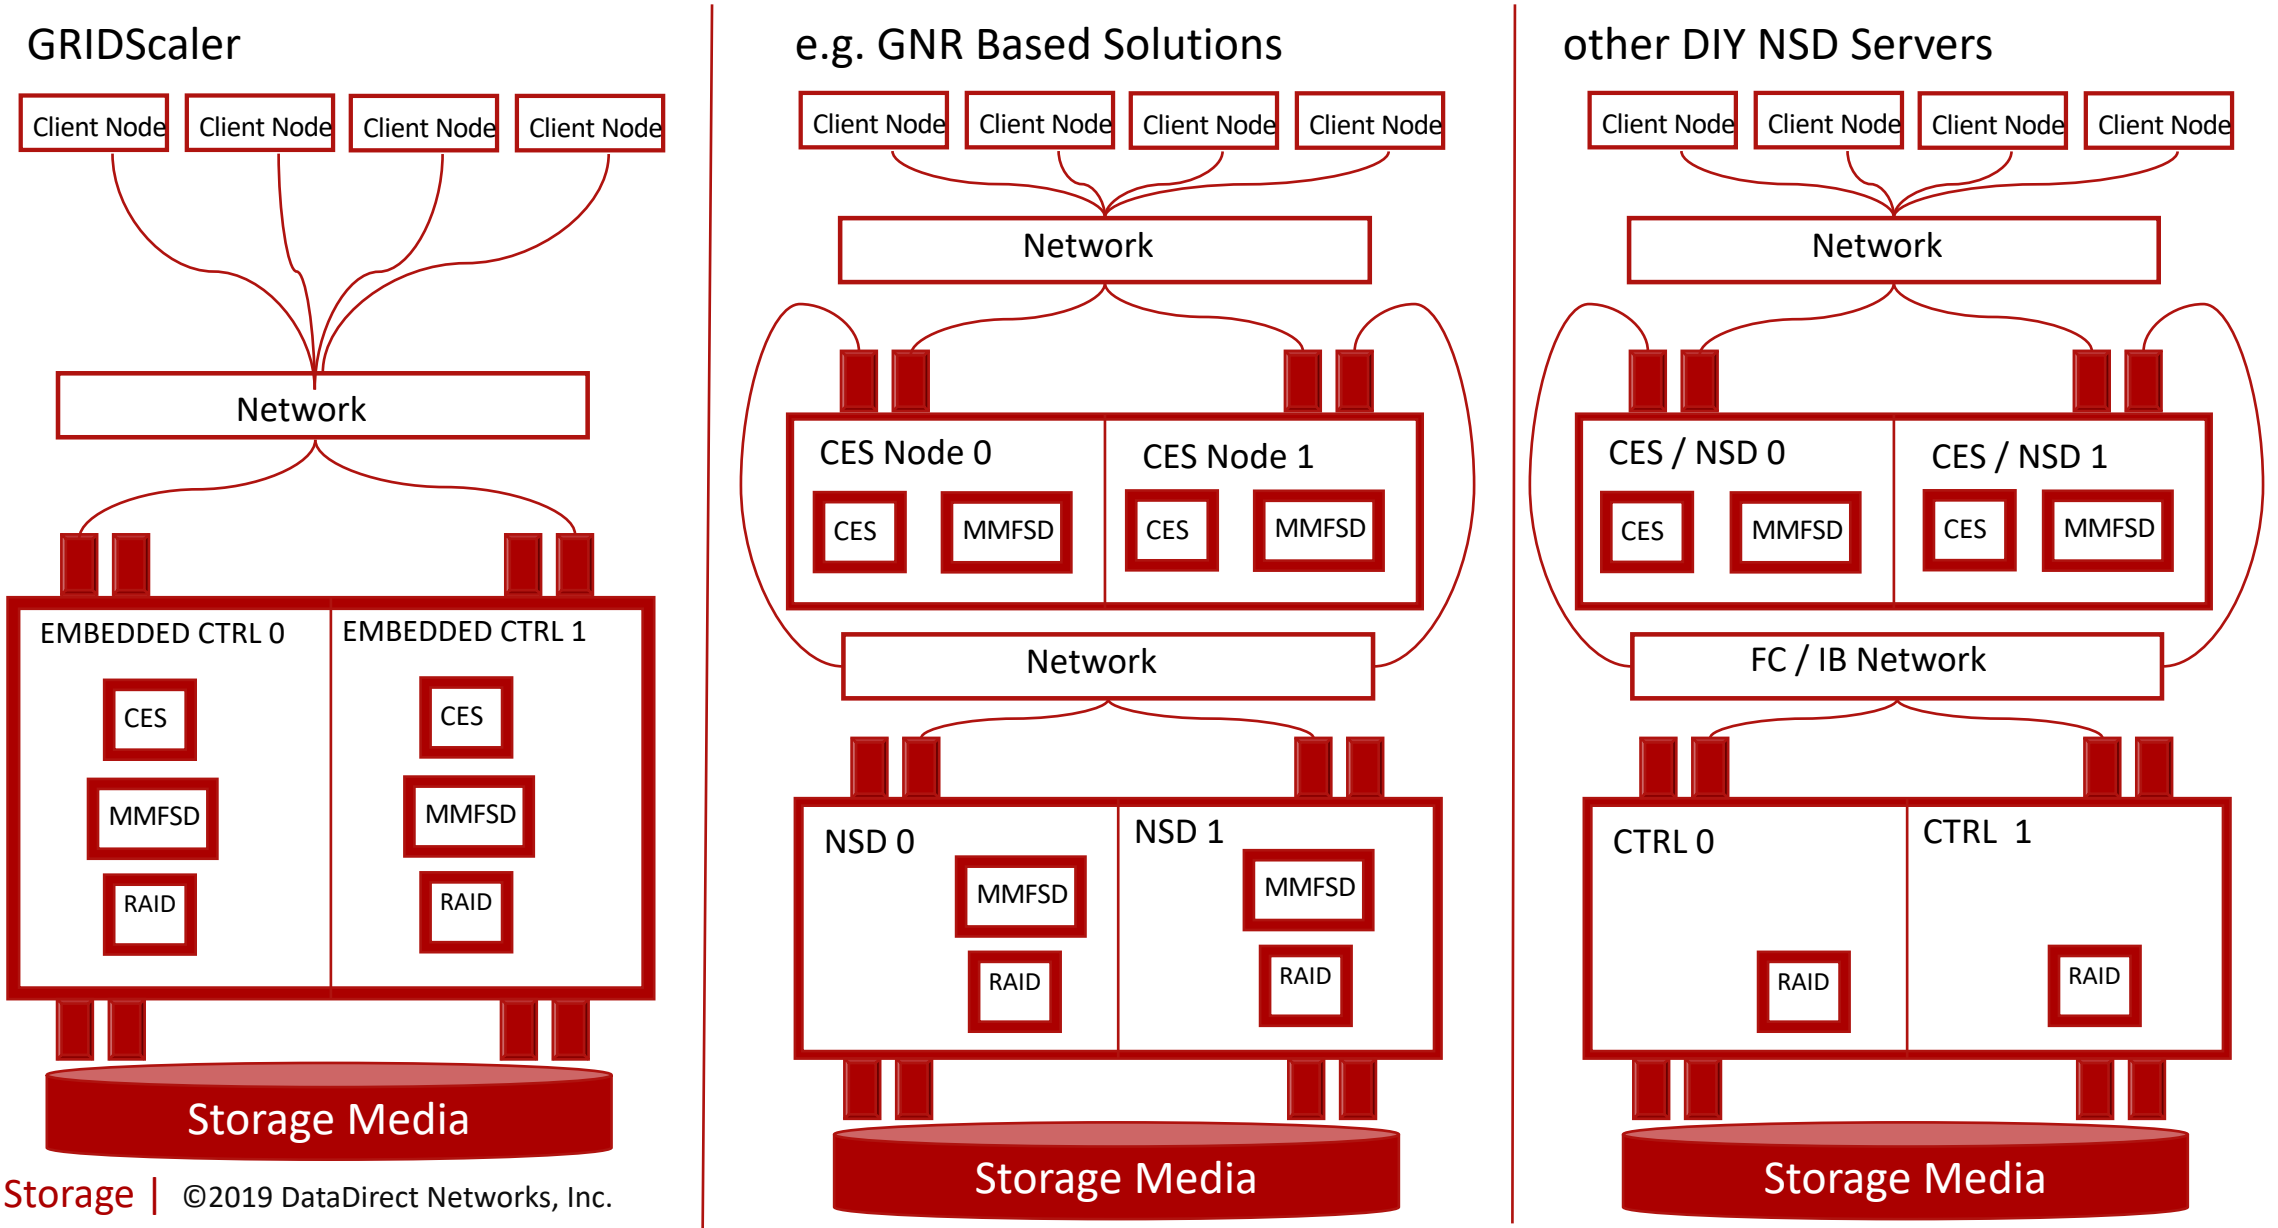

SFA NUMA awareness

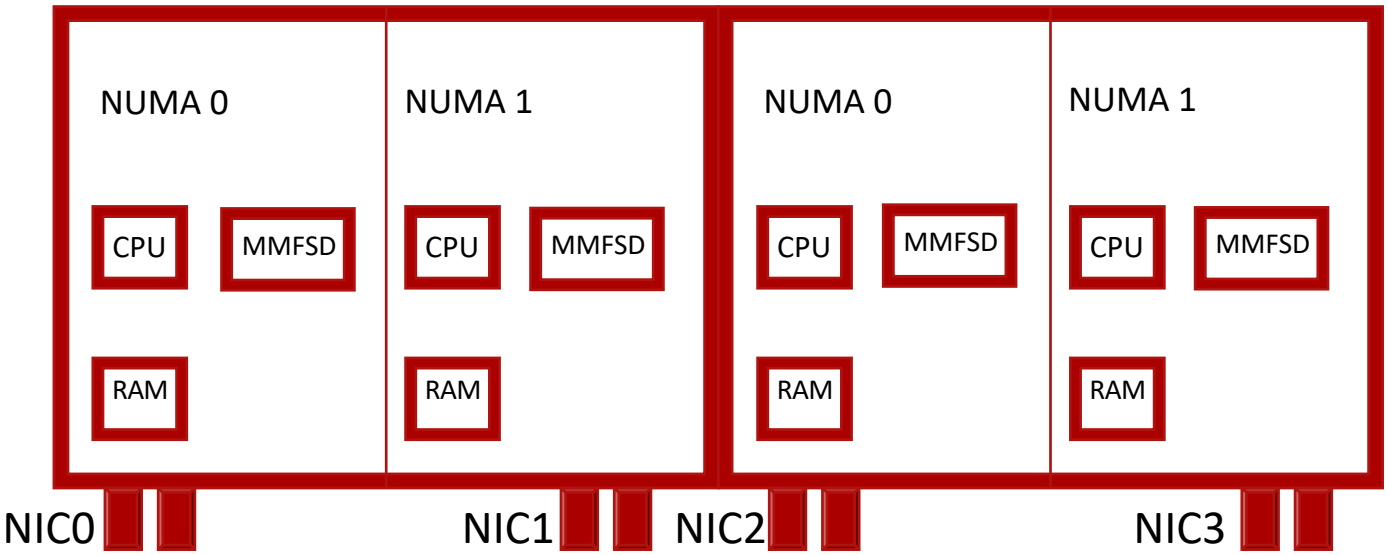

The system is perfectly balanced across numa nodes, which allows affinitizing of mmfsd threads to memory, core and network for lowest latency and consistent scaling

18k results with GRIDScaler V5 and SFAOS 11.5

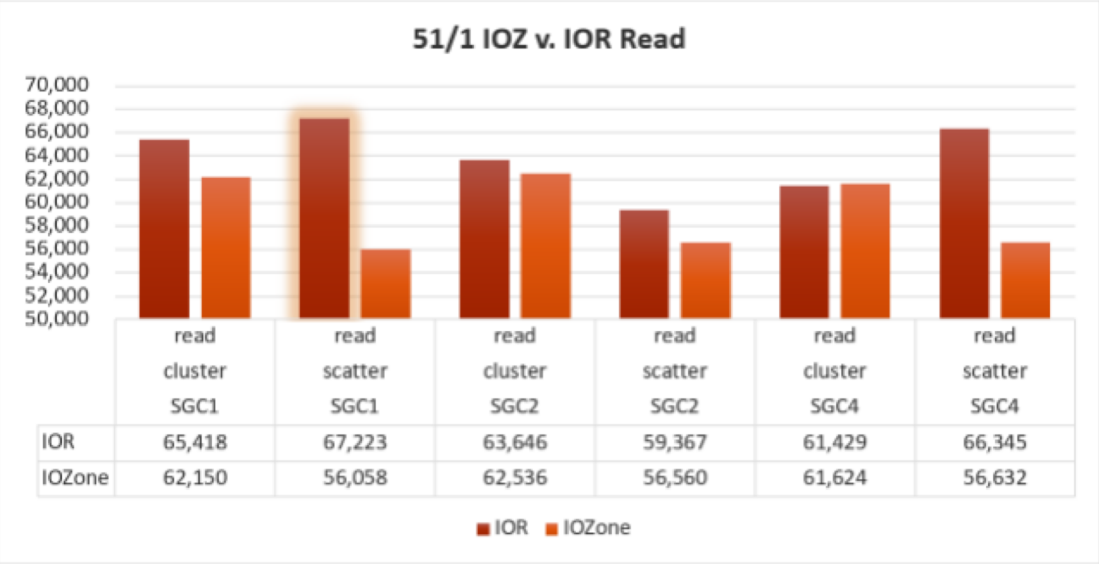

	read cluster SGC1	read scatter SGC1	read cluster SGC2	read scatter SGC2	read cluster SGC4	read scatter SGC4
IOR	65,418	67,223	63,646	59,367	61,429	66,345
IOZone	62,150	56,058	62,536	56,560	61,624	56,632

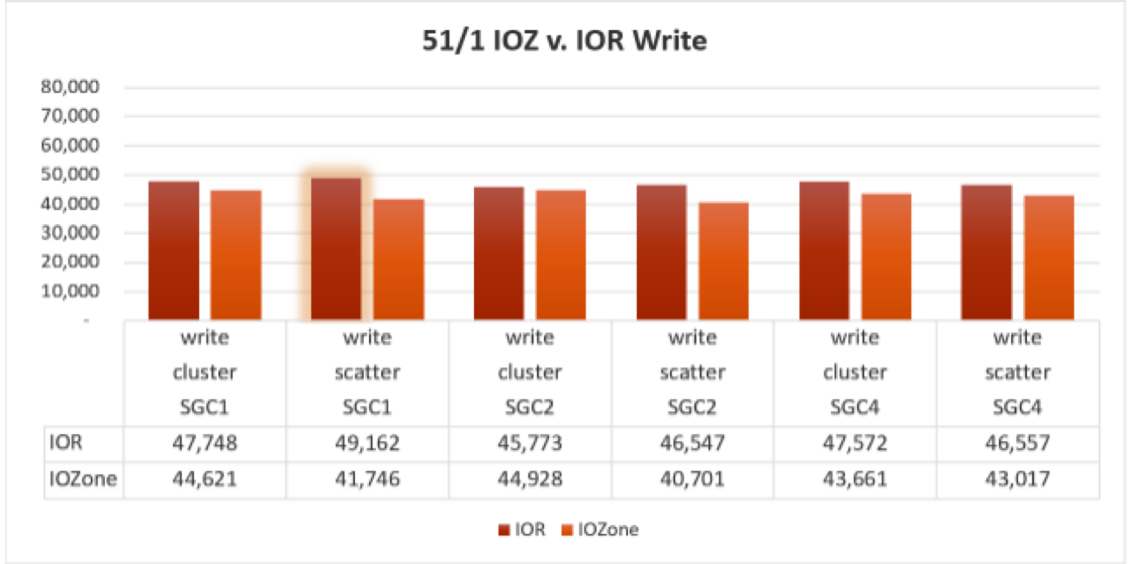

	write cluster SGC1	write scatter SGC1	write cluster SGC2	write scatter SGC2	write cluster SGC4	write scatter SGC4
IOR	47,748	49,162	45,773	46,547	47,572	46,557
IOZone	44,621	41,746	44,928	40,701	43,661	43,017

System Setup is a single SFA18K with 408 HDD's running SFAOS 11.5, with 8 Pools, each 51/1 RAID 6 with 2MB chunk size

We tested with 16 clients connected via single port EDR cable, data set size was >10x of all combined caches in each run

This translates in ~150 MB/sec per individual HDD for full Block random Read/writes in scatter mode

~150MB/sec/ HDD

DIO Random 4k writes into a 100GB files

```
/usr/lpp/mmfs/samples/perf/gpfsperf write rand /target/sven-100g
  recSize 4K nBytes 100G fileSize 100G
  nProcesses 1 nThreadsPerProcess 1
  file cache flushed before test
  using direct I/O
  offsets accessed will cycle through the same file segment
  not using shared memory buffer
  not releasing byte-range token after open
  no fsync at end of test
    Data rate was 34659.88 Kbytes/sec, Op Rate was 8461.89 Ops/sec, Avg Latency
was 0.118 milliseconds, thread utilization 1.000, bytesTransferred 1039802368
```

DIO Random 4k reads from a 1TB files (exceeds all cache by >4x)

```
/usr/lpp/mmfs/samples/perf/gpfsperf read rand /target/testfile-1t
  recSize 4K nBytes 1024G fileSize 1024G
  nProcesses 1 nThreadsPerProcess 1
  file cache flushed before test
  using direct I/O
  offsets accessed will cycle through the same file segment
  not using shared memory buffer
  not releasing byte-range token after open
    Data rate was 27982.49 Kbytes/sec, Op Rate was 6831.66 Ops/sec, Avg Latency
was 0.146 milliseconds, thread utilization 1.000, bytesTransferred 279830528
```

DDN® DataFlow

Protect, Vault, Move, Synchronize Data at Scale

DDN DATAFLOW SOLUTIONS FOR DATA MANAGEMENT AT SCALE

Backup, archive and move data between any storage

- Scalable, high-performance, flexible architecture
- Interface with a wide range of data platforms
- FastScan technology for rapid filesystem scanning
- Compression, encryption, checksum, versioning
- Extensive access control, activity logging, auditing

- True self-service for users at all levels of expertise
- Intuitive GUIs for administrator and user control
- Comprehensive APIs and CLI for integration

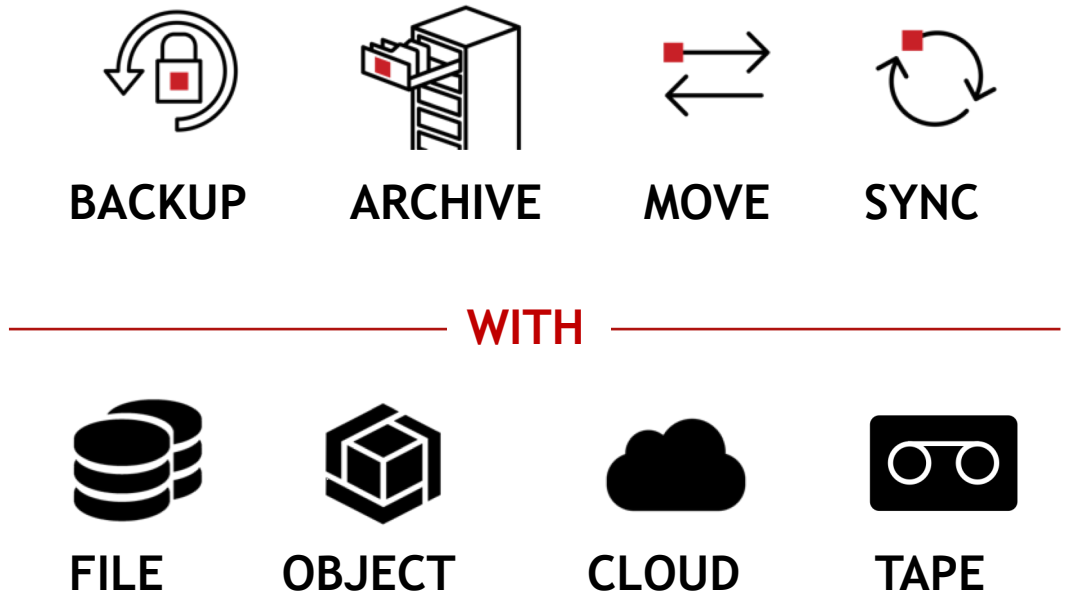

MIGRATE ALL YOUR DATA RAPIDLY, EASILY, AND WITH CONFIDENCE

DDN® DataFlow

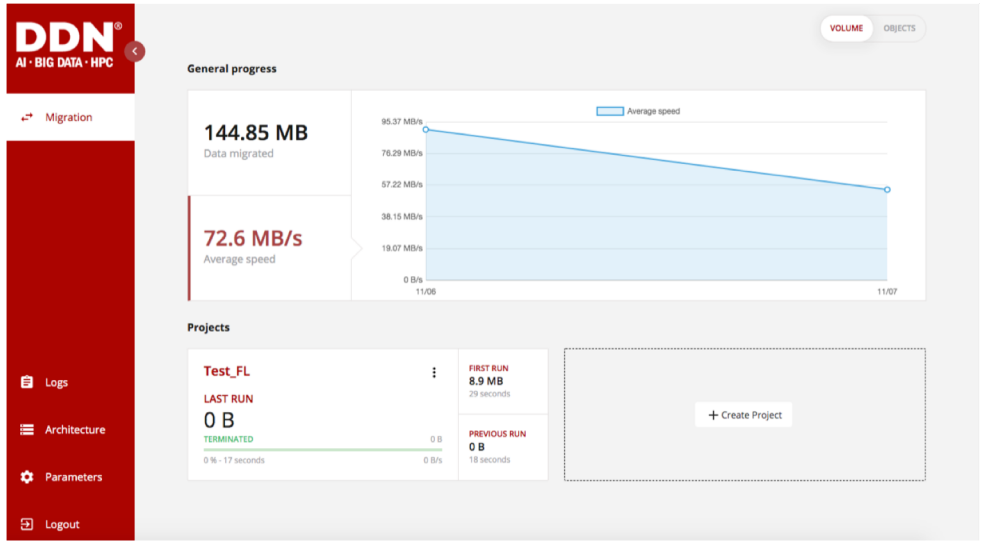

- At-scale data migration solution for massive data sets
- Move your data, maintain permissions and ACLs
- Migrate your data safely, end-to-end data integrity checks
- Automated migration through scheduled jobs
- Follow migration process easily through an intuitive GUI

DDN DATAFLOW MIGRATION CENTRALIZED MANAGEMENT

Intuitive user interfaces for effortless productivity

The administrator console provides single pane of access for complete system configuration, workflow definition and process monitoring.

Historic and real time information of the migration tasks is available enabling customer to easily follow the migration process at all time.

Comprehensive CLI, web services and a C++ API are also available for automation and integration.

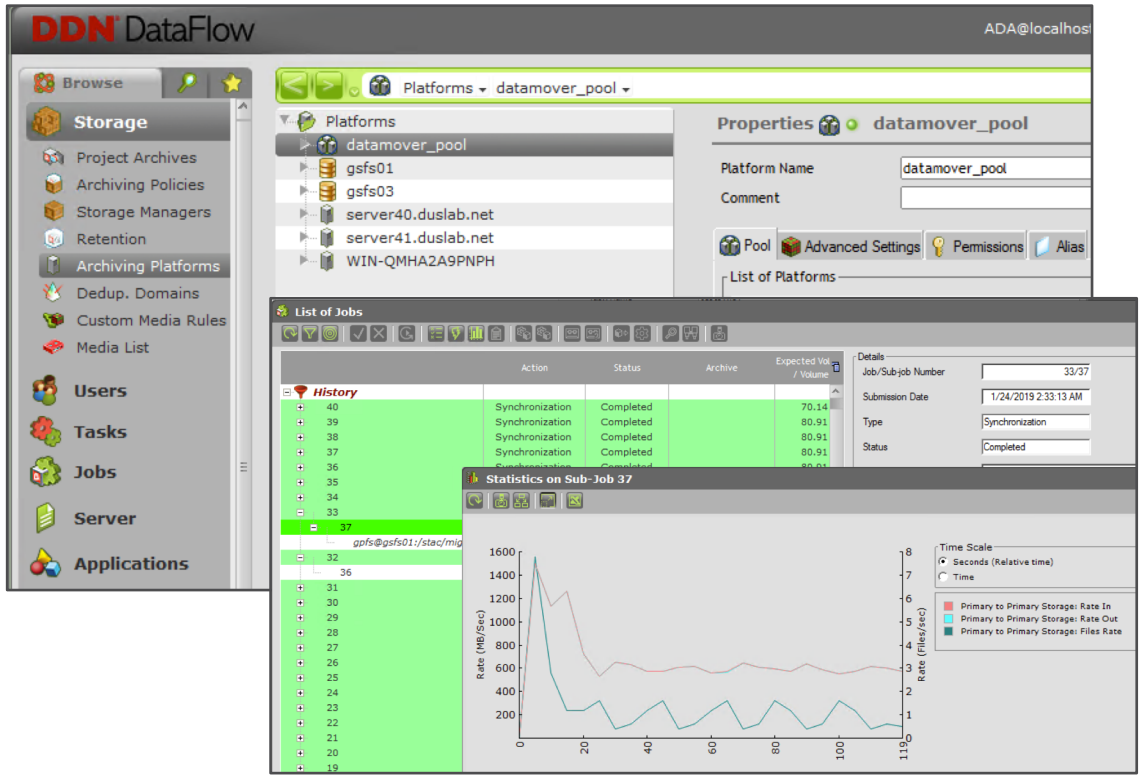

DDN DATAFLOW TURNKEY AT-SCALE DATA MIGRATION SOLUTIONS

EXPERT DATA MIGRATION SOLUTIONS FROM DDN

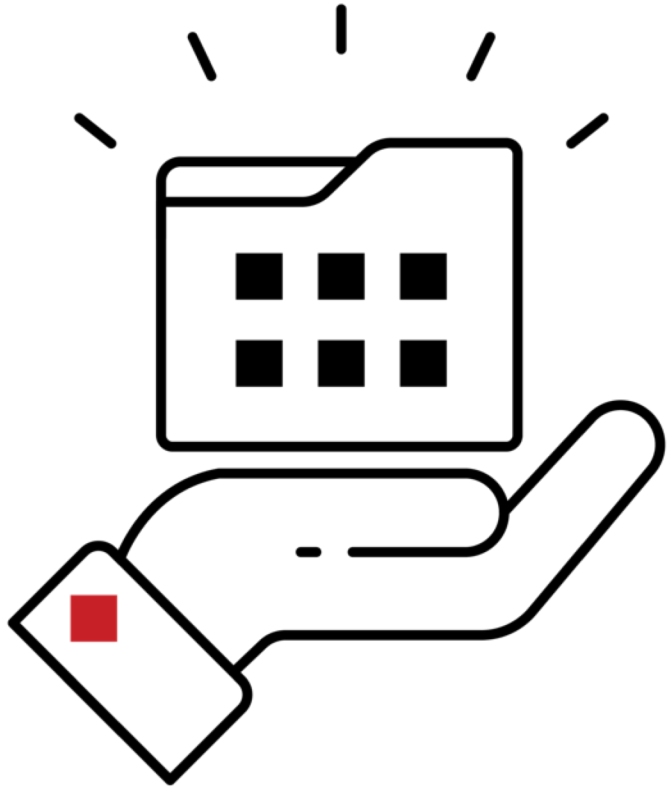

Pre-configured appliances for data migration that are easy to deploy and integrate within any IT infrastructure and storage environment.

Experienced DDN engineers handle planning, execution and validation of data migration operations at any scale.

DDN professional services are available worldwide onsite and remotely.

DDN.COM/DATAFLOW

Infinite Memory Engine

Limitless performance for your most demanding workloads

Modern Workloads mean bigger pressures for filesystems

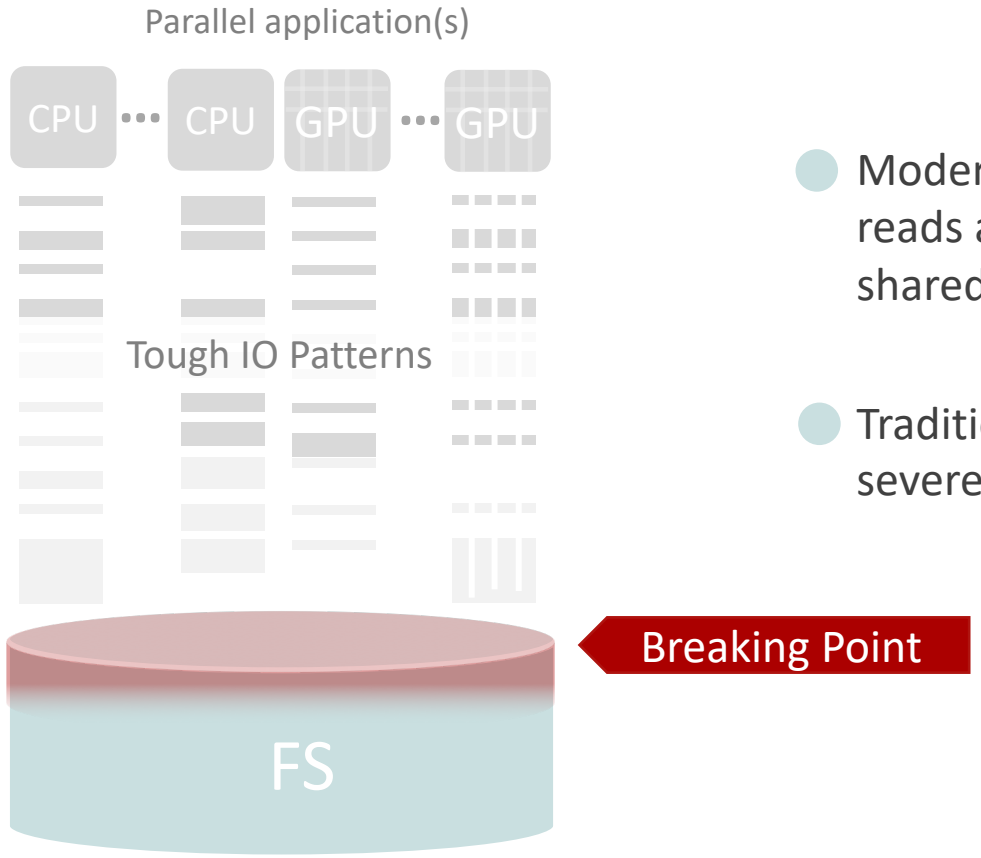

- Modern Workload IO patterns are increasingly mixed and tough: reads and writes, random and sequential, high thread counts, shared file access
- Traditional **Thick File system SW layers** and **fixed data layout** severely restrict performance for tough workloads – even with SSDs

What about NVMe over Fabrics?

That solves all latency problems, right?

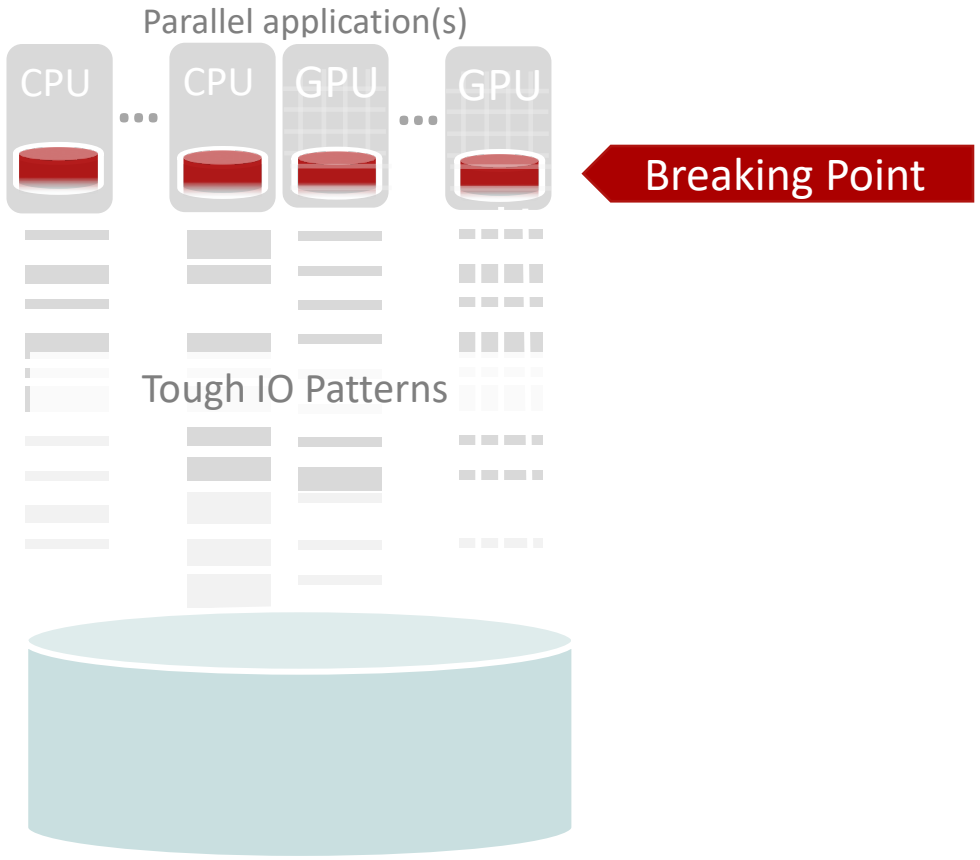

- NVMeoF solves a problem for block, but not for file access. It just moves the bottleneck
- Regardless of the method for providing a block device there is still latency in the filesystem layers
- Your applications only care about filesystem latencies and throughputs

IME's Optimal Data Path accelerates your I/O

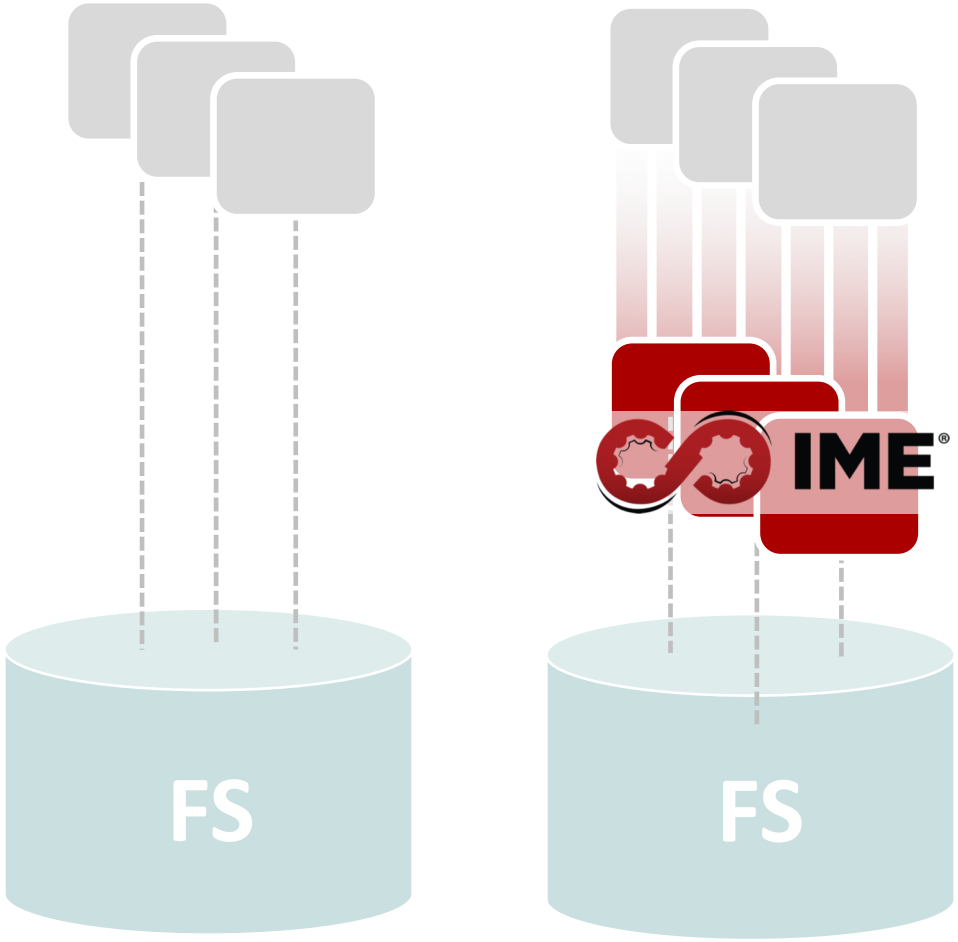

- IME forms a transparent, scalable cache which delivers unprecedented performance to applications
- Zero Application modifications are needed for IME to unleash the power of your next generation workloads
- IME dramatically accelerates random read, random write, shared file, high concurrency and streaming workloads

IME enables new levels of filesystem performance

- Parallel File systems can exhibit extremely poor performance for shared file IO due to internal lock management as a result of managing files in large lock units

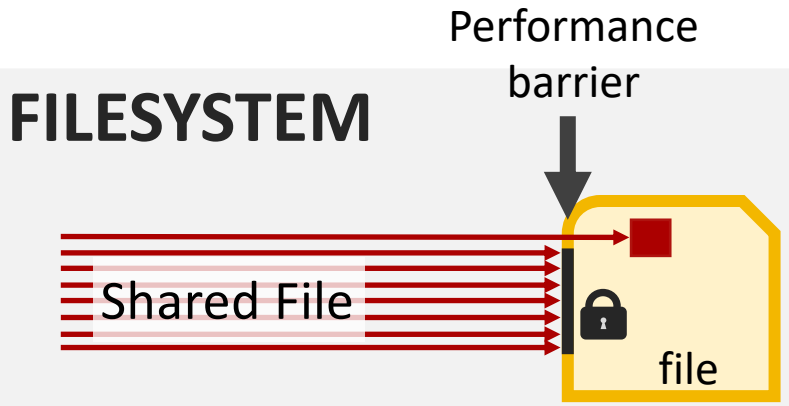

- IME eliminates contention by managing IO fragments directly, and coalescing IO's prior to flushing to the parallel file system

IME | SW-Defined Performance and Resilience

- Developed from scratch for Flash, IME delivers a new highly differentiated feature set, faster rebuilds, selectable resilience and adaptive I/O

Adaptive IO

IME Clients adapt IO rates to the server layer according to load eliminating traditional IO system slowdowns

Dial-in Resilience

Erasure coding levels are not dictated by storage system setup, but dynamically set on a per file/per client basis

Lightning Rebuilds

Fully Declustered distributed Data rebuilds allow for rebuild rates in excess of 250GB/minute

Open Monitoring with IME and Grafana

- Monitor Job Activity
- Monitor IME capacity, throughput and IOPs
- Monitor activity of the Backing filesystems
- Job activity Overview

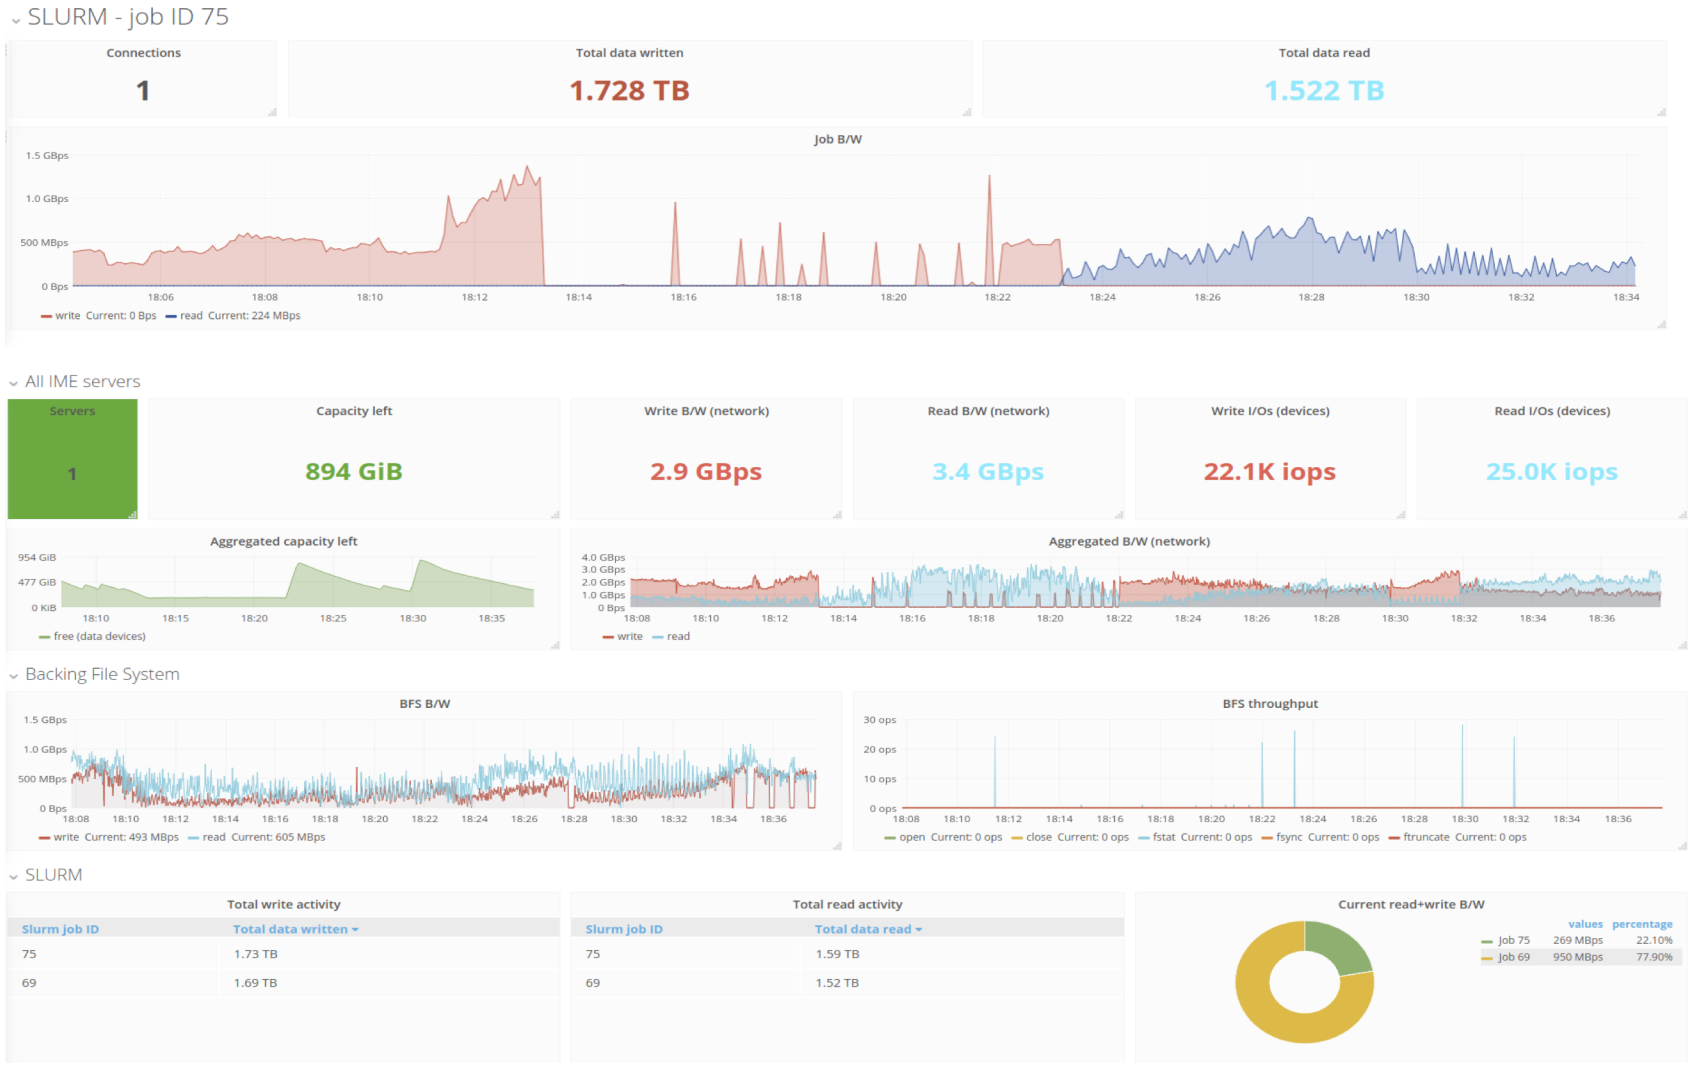

IME efficiently delivers EXAScale Performance

- Real world implementation of around 2 racks of IME
- ~1PB Flash
- Lustre Backing Filesystem
- measured 1.2 TB/s
  - Both File per Process AND
  - Single Shared File

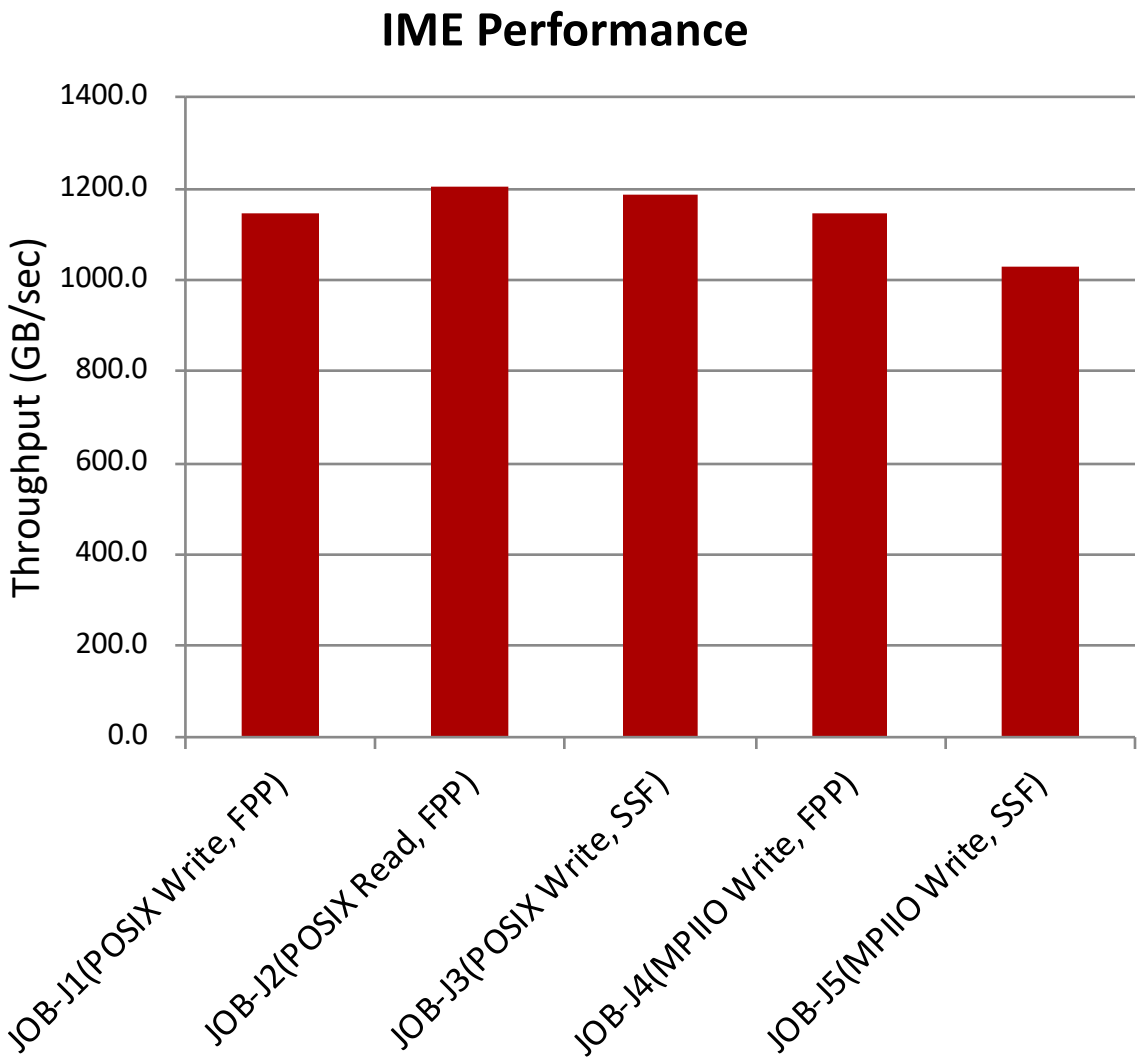

IME 1.3.1 – IME on GCP and AWS

- IME 1.3.1 for AWS is deployment ready
  - IME performance close to the Instance performance
    - Limited only by Network bandwidth
  - IME sustains functionalities with the FS On cloud and On premise
    - Full IME functionalities with high latency configuration
    - Sustained BW values

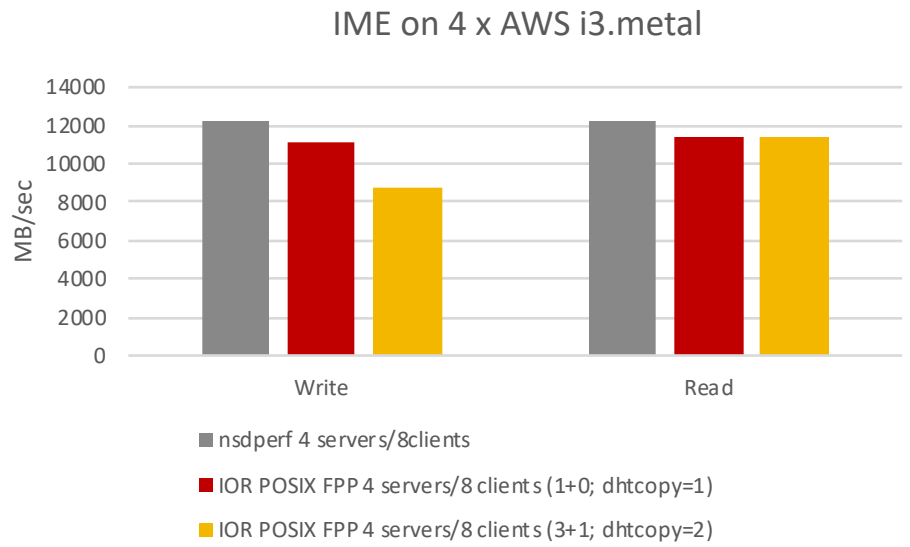

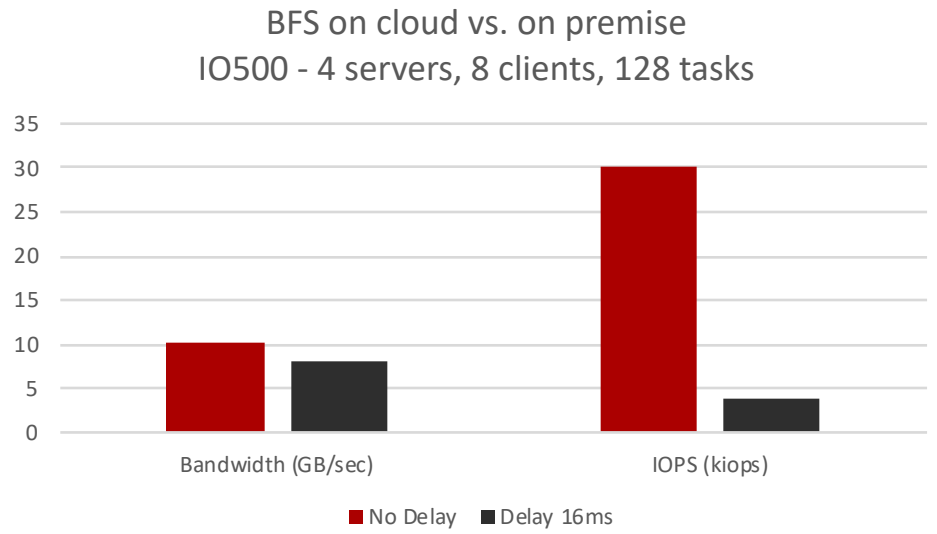

IME – Burst buffer for NFS

- Brings scale-out Flash native performance to NFS access
- Shield NFS server from ”tough" IO
- Increase IO throughput from NFS hardware
- Zero application changes - replace NFS mount by IME mount

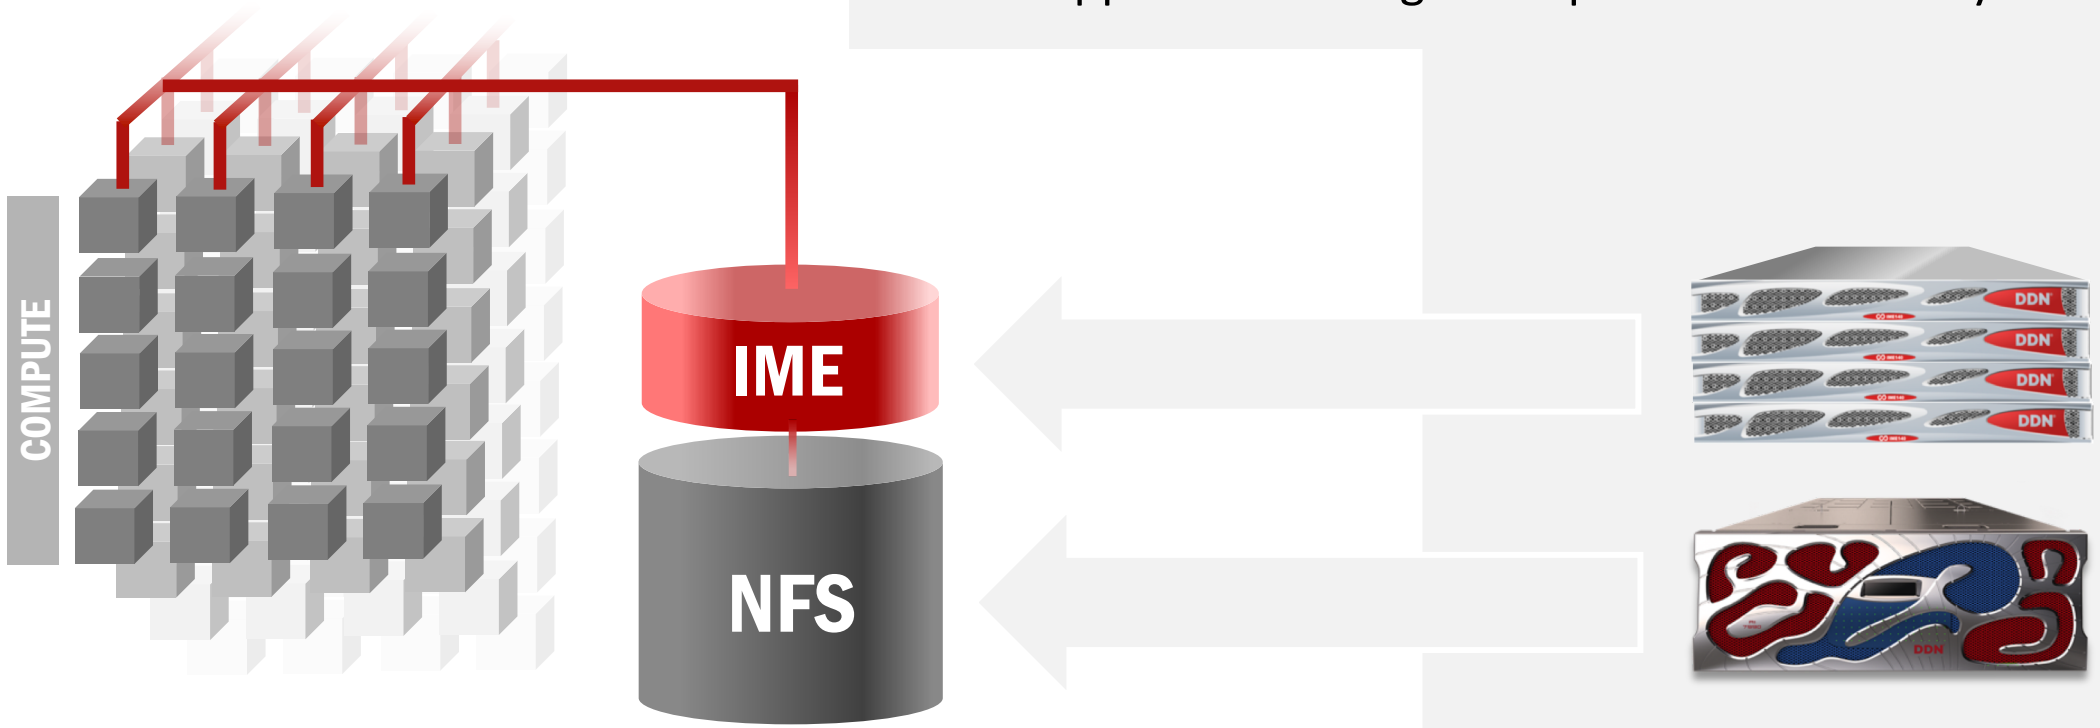

IME – Burst buffer for NFS

IME with NFS

- Brings scale out Flash native performance to NFS Systems
- Removes complexity associated with Parallel Filesystems
- Shield NFS server for "bad" IO
- Increase IO throughput on top of NFS hardware
- No application changes - replace NFS mount by IME mount

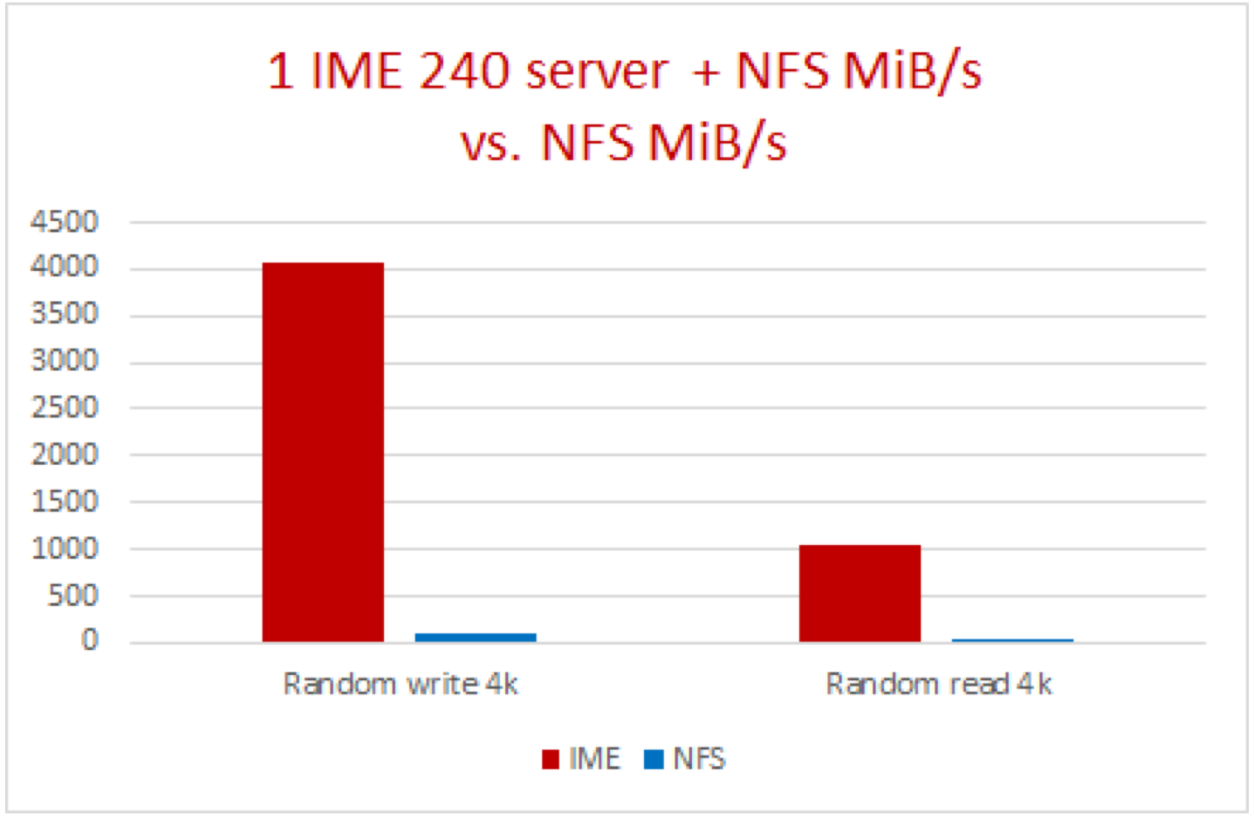

Thank You!

Keep in touch with us.

sales@ddn.com

@ddn_limitless

company/datadirect-networks

9351 Deering Avenue
Chatsworth, CA 91311

1.800.837.2298
1.818.700.4000