

The Path to Finding Business Value in Data



Molly Rector

**DDN CMO, EVP Product Management &
WW Marketing**

SNIA, Vice-chairman, Analytics & Big Data Committee

Active Archive Alliance, Board Member



2 **DDN** | Enabling Enterprise Big Data Initiatives


Financial Services

Powering 40% of Leading
Global Investment Banks



Manufacturing

Powering 30% of the World's
Top Aero & Auto Manufacturers



Web, Cloud & Telco

Powering the Largest Web Scale
Global Service Providers



Media

Powering over 600
HD Workflow Customers



Supercomputing

Powering Over 2/3 of the Top100®



Life Sciences

Powering Over 1/3 of the
Top Sequencing Centers



DDN | Powering the Largest Global Enterprises

Financial Services

Powering 40% of Leading
Global Investment Banks



Manufacturing

Powering 30% of the World's
Top Aero & Auto Manufacturers



Web, Cloud & Telco

Powering the Largest Web Scale
Global Service Providers



Media

Powering over 600
HD Workflow Customers



Supercomputing

Powering Over 2/3 of the Top100®



Life Sciences

Powering Over 1/3 of the
Top Sequencing Centers



DDN | High Performance Big Data Application Optimization

Financial Services

Powering 40% of Leading
Global Investment Banks

Portfolio Risk Management
Complex Back Testing
Same Day Trading

Manufacturing

Powering 30% of the World's
Top Aero & Auto Manufacturers

Engineering Modeling
Crash Simulation
Chip Design

Web, Cloud & Telco

Powering the Largest Web Scale
Global Service Providers

IOT Content Capture
Data Lake Optimization
Efficient Content Distribution

Media

Powering over 600
HD Workflow Customers

Ingest, Edit, Animate
Content Collaboration
Active Archive

Supercomputing

Powering Over 2/3 of the Top100®

Application Optimization
Highest Performance Mixed I/O
Efficient Storage Condos

Life Sciences

Powering Over 1/3 of the
Top Sequencing Centers

Genomic Sequencing
Clinical & Epidemic Predictions
Drug Discovery

It's About Getting to the Data

- **Commercial users leverage HPC to derive higher financial returns from data**
 - Innovation
 - Competitiveness
 - Time to Market
 - Quality
 - Profitability
- **If the problem is big data, then the solution needs to focus on getting to lots of data very fast**
- **Storage, Networking and Applications are more likely to be bottlenecks than compute**

DDN Customers

Leveraging Big Data to Deliver Results

VIDEO

DataDirectTM
N E T W O R K S

PIONEER SERIES

**ACCELERATE: GENOMICS RESEARCH**

https://www.youtube.com/watch?v=9aLmnqWv9AI&index=2&list=PLohLhg8SIVmY_F-6okdiu_g2DxP7BuAcg

The Big Data Market is Exploding!!!

2014-2019 CAGR

- 9%: HPC Storage Market
- 29.9%: High Performance Data Analytics Segments

Source:

Growing Opportunities

- ✓ Hybrid Storage
- ✓ Increased SSD software
- ✓ Active and deep archive
- ✓ Flash / Cache

Users moving
AWAY from
traditional IT
platforms for high
performance

apps

Identify the Right Tools for the Job



- Traditional IT tools are not built for the mix of media types nor combination of speed and scale needed to solve big data problems
- In storage, typically, 2 to 20 traditional IT systems can be replaced with 1 system designed for Big Data
 - Huge difference in moving parts
 - Better footprint and energy profile
 - Still has room to grow

9 Requirements of the Big Data Environment



✓ **Sustained Performance (Mixed IO and at Scale)**



✓ **Storage and Performance Cost Density**



✓ **Management Simplicity at Scale**



✓ **Integrated Tiering**



✓ **Embedded Multi Site Collaboration**

DDN | Comprehensive Product Portfolio

Powerful, Unified, Delivering Results...Faster



Compute & Memory



Fast Data
(NVM & SSD)



1000x IO Acceleration

100x File System Acceleration

10x Application Acceleration



Persistent Data
(SSD & Disk Hybrid)




SSD, SAS, SATA

Embedded F/S and Apps

Mixed Workloads



Distributed Data
(Object)



Collaboration and Global
Distribution, Private and
Public Cloud

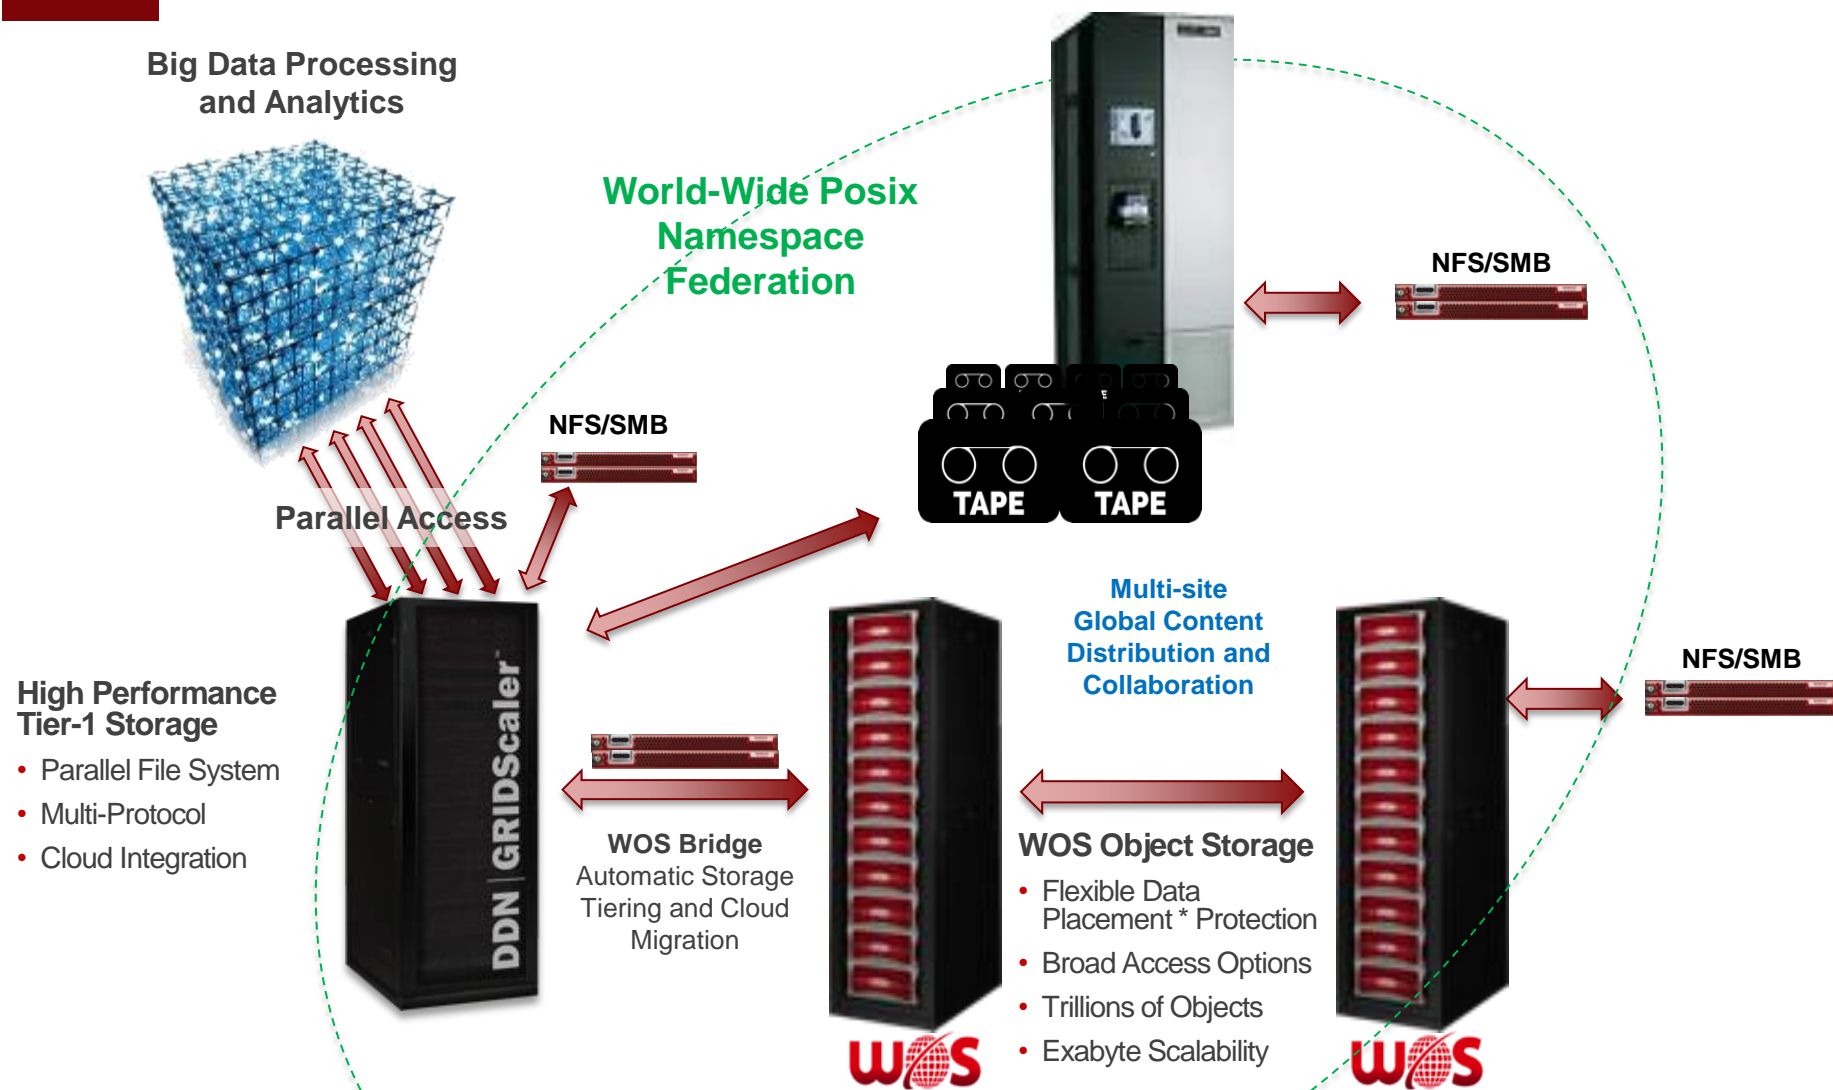
Massive Scale & Low Latency

Cloud & Archive

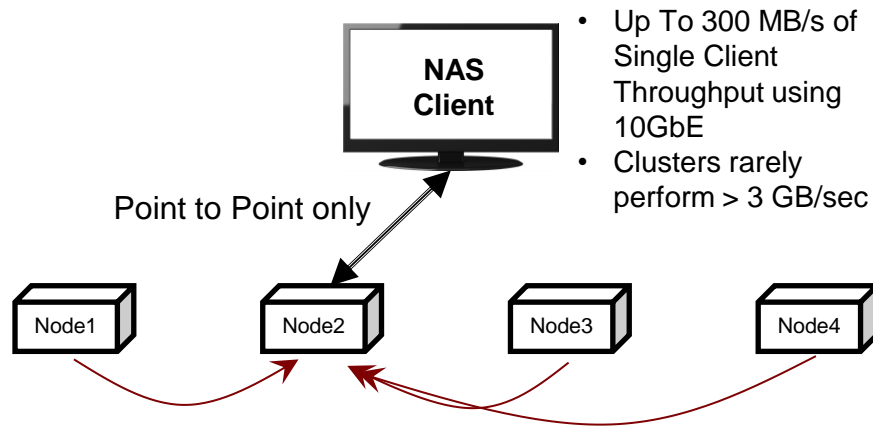


DDN | World-Wide Namespace Integration

Collaborate, Distribute and Federate

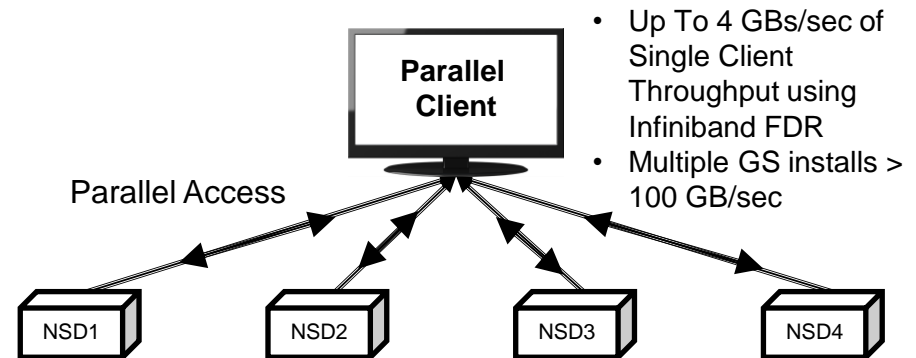


Scale-out NAS



- Point-to-Point Technology
 - Client cannot read from multiple servers in parallel
 - Performance limited by single Node capability
- Not designed for Concurrent Access of files/directories
- No support for IB RDMA Access

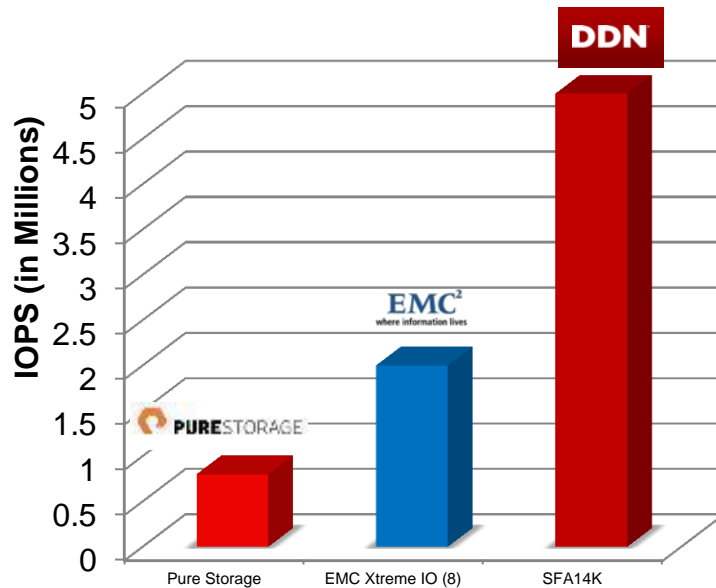
Performance File Storage



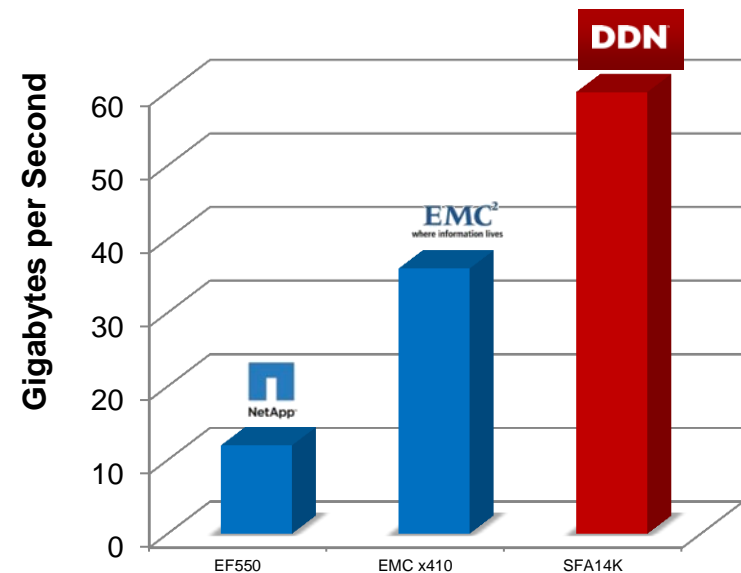
- ▶ Supports NAS as well as Parallel File System protocols
 - Designed for Massive Parallelism
 - Accelerate application performance by striping data across multiple enclosures and accessing them in parallel
- ▶ Support IB RDMA – the leading low latency protocol

#2 – Small File, Mixed IO, Performance & Scale

Transactional Performance Appliance IOPS



Streaming Performance Appliance Bandwidth



Delivering Optimal Solution Platforms



When a New Node is added, the File System has to be restriped/rebalanced to accommodate additional nodes



- This data copy can take weeks and can choke the cluster
- It kills performance

“In high load video processing, restriper jobs may cause contention related latency leading to dropped frames.”

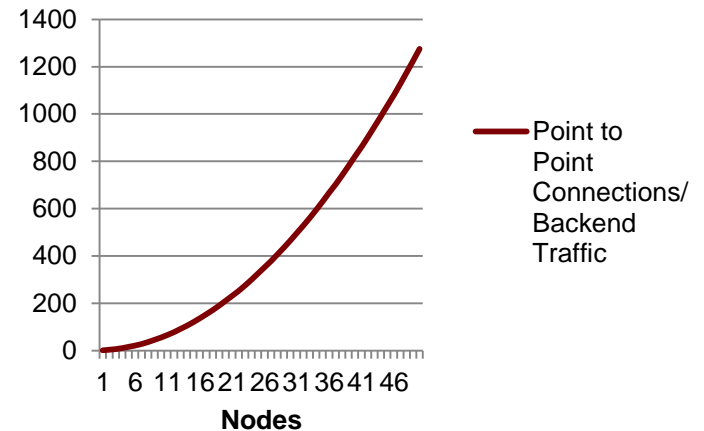
- In other words – performance penalty is severe enough to drop incoming data

Potential Technology Limitations

• As the Cluster grows

- The cluster sees an exponentially increase in metadata traffic and node communication/processing
- As cluster grows, the data service responsibility per node increases a) as a node has conflicting requests – it's own and the cluster's

Exponential Increase in Traffic with Increasing Nodes



Performance
Drop

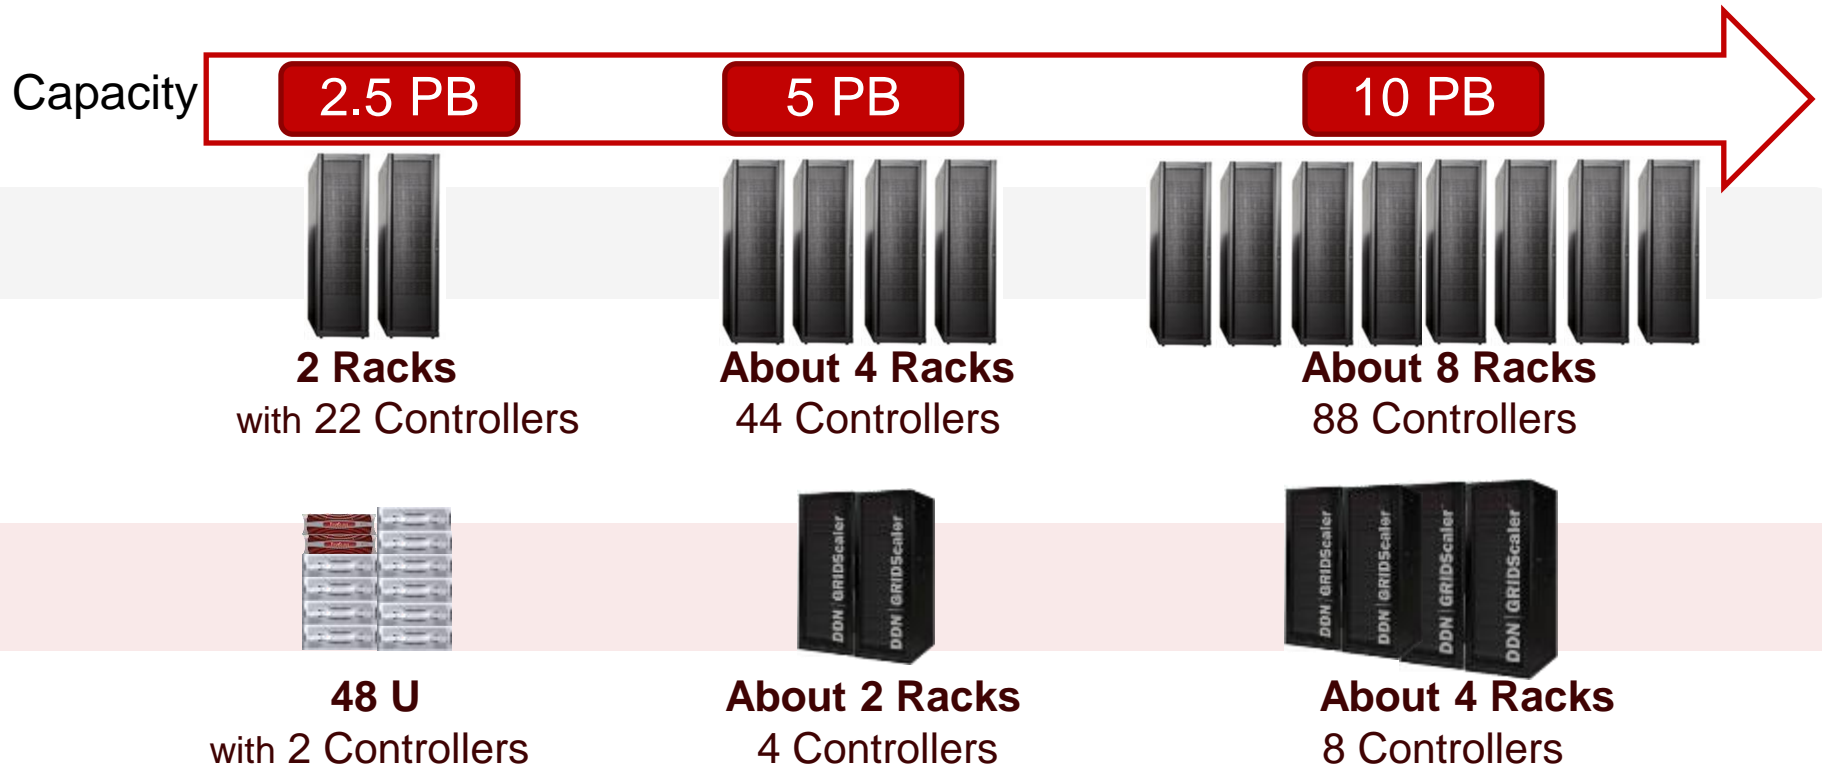
Latency Increase

DATA

PARITY
METADATA
REBUILDS!

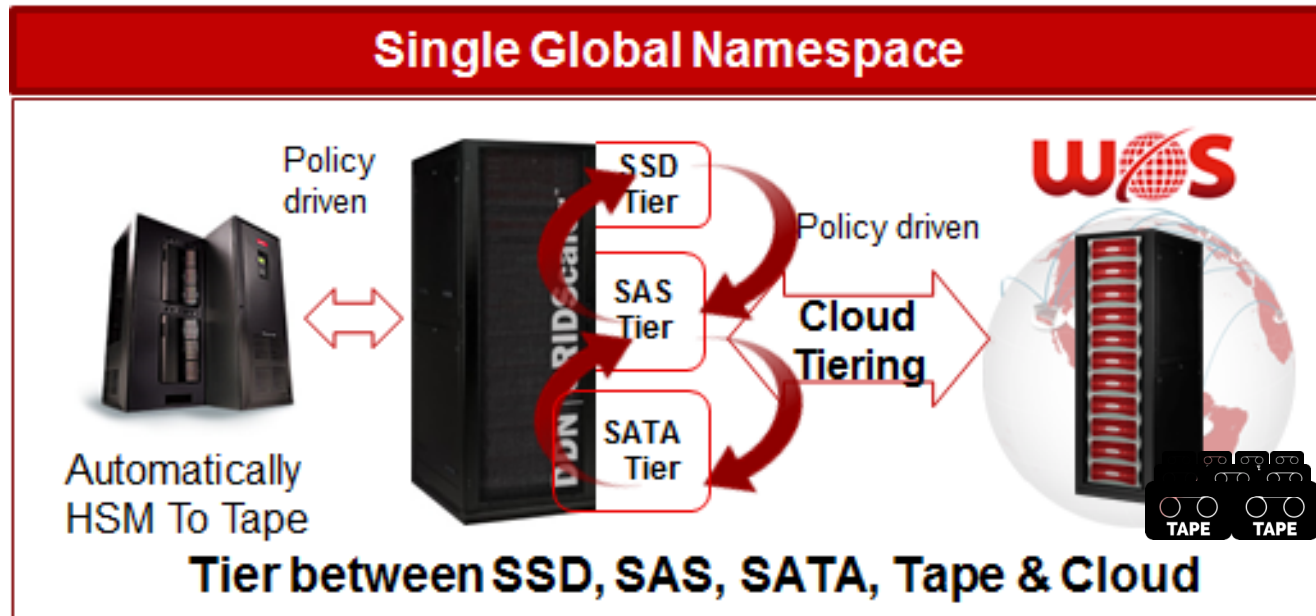
#4 - Storage and Performance Density

- ✓ Simplified Management
- ✓ Lowered space and power utilization



Comparison using DDN: GridScaler 12KX with SS8460 Enclosures, Isilon: NL400 with N+2 parity – 144TBs in 4U


#5 – End-to-end Tiering



Parallel NAS

- All Tiers under single global namespace
- All data easily accessible
- Tier to a Private Cloud

- ▶ Federation of namespaces
- ▶ Multi-site ingest, distribution & collaboration
- ▶ Publish data at any location and access from variety of protocols



Full Spectrum Object Storage Solution



Security

- ▶ End-to-end network security

Accessibility

- ▶ Wide choice of access protocols, multi-site distribution with federated namespace

Efficiency

- ▶ Low overhead, no bandwidth abuse, lowest management cost

Reliability

- ▶ Flexible policy driven data protection and self-healing architecture

Performance

- ▶ Highest performance for throughput, IOPS, and low latency

Scalability

- ▶ Seamlessly scale performance, number of objects, capacity and distributed sites

High-Performance Converged Media Storage Platform

Clustered NAS • End-to-End Data Management

Turnkey Scalable Storage Solution

Animate

Render

Edit

Transcode

Collaborate

Active Archive



Use Cases

- Public/Private Cloud Storage
 - Unlimited scale, Predictable performance, Global access
- Active Archives
 - Storage efficiency, Data durability, No forklift upgrades
- Content Distribution and Collaboration
 - Policy-based data movement, Network efficiency, Global access
- Disaster Recovery
 - Multi-Site, Active-Active (-Active...)
- Video Streaming
 - Low latency, predictable performance



RENT



BUY



Build your Big Data infrastructure in house
OR
Outsource some or all of it
Compute + Storage + Network + Expertise



- HPC is a competitive weapon for demanding Enterprise workflows to turn data into value
- The problem is big data - the solution needs to focus on getting to lots of data very fast
- Don't bring a knife to a gun fight – use the right tool for the job



Thank You!

Keep in touch with us



sales@ddn.com



2929 Patrick Henry Drive
Santa Clara, CA 95054



[@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)