March 2013

DataDirect™ NETWORKS
INFORMATION IN MOTION™

DataDirect Networks Update

SOS17 Conference – Jekyll Island, Georgia

John Josephakis

SVP of HPC

DDN HPC Investments Today, The Enterprise Data Center of Tomorrow

DataDirect NETWORKS
Information in Motion

DDN's $100M Commitment To The Future of HPC

- I/O Acceleration For Highly Concurrent Systems
- Exploitation of Next Generation Non-Volatile Memory
- Ultra-Low Power and Infrastructure Efficiency

DDN Engineering Investment

- Within the next 8 years, DDN will spend over $500M in Engineering dollars.
- We have apportioned a significant amount of this budget for Exascale R&D & to solve the toughest I/O challenges.
- DDN has uniquely mastered the art of commercializing @scale HPC storage tech.

©2012 DataDirect Networks. All Rights Reserved.

DDN & DoE FastForward

DataDirect NETWORKS
Information in Motion

DDN Is A Key Partner to the DoE/NNSA Exascale I/O R&D Program

- Co-Development of Exascale I/O Layers
- Lustre®, Burst Buffer, Object Store
- $M Long-Range Development Effort
- 100% Open Source Contribution
- Complete Alignment with DDN Exascale Strategy

©2012 DataDirect Networks. All Rights Reserved.

DDN's Focus is HPC.

DataDirect NETWORKS
Information in Motion

DDN Powers the Fastest HPC Systems in the World

The Leader on 5 Continents and >66 of the Top100
Delivering More GB/s Than All Others *Combined*

TOP500® SUPERCOMPUTER SITES

>170 of the Top500
Growing 20% Annually

DDN provides more total storage bandwidth to the Top 500 than all other storage vendors combined

©2012 DataDirect Networks. All Rights Reserved.

An Amazing Step Forward…

DataDirect NETWORKS
Information in Motion

OAK RIDGE National Laboratory

Spider-Atlas

DDN and ORNL Have Partnered To Build The World's Fastest Storage System; Supporting Titan

System Performance: 1TB/s+
Capacity: 40.3PB (raw)
File System: Lustre®
I/O Platform: 36 x DDN SFA12K-40
Media: 20,160 HDDs

Fun Fact:
ORNL Titan Is Designed With The I/O Bandwidth Equivalent To 80,000x the Amount of Tweets and Tweet Metadata per Second from Twitter.

©2012 DataDirect Networks. All Rights Reserved.

The Opportunity is Growing

“’big data’ has long been an important part of the HPC market, but recent technology advances have given data-intensive computing much higher potential as a horizontal market.”

IDC
2012

©2012 DataDirect Networks. All Rights Reserved.

DDN | World-Leading Deployments

DataDirect NETWORKS
INFORMATION IN MOTION

HPC & Big Data Analysis

- Wellcome Trust Sanger Institute
- AIRBUS
- SUZUKI
- TOTAL
- CEA
- SAS
- STANFORD UNIVERSITY
- OAK RIDGE National Laboratory

Cloud & Web Infrastructure

- CDNetworks
- Level 3 COMMUNICATIONS
- ebay
- IBM
- gettyimages
- SAMSUNG
- FAMILYSEARCH WHERE GENERATIONS MEET
- yousendit

Professional Media

- UNIVERSAL
- CBS SPORTS NETWORK
- SONY
- SPEED CHANNEL
- MEDIASET
- qvc
- deluxe

Security

- Department of Defense
- Public Transport Authority
- Sheriff Los Angeles County
- Raytheon
- BAE SYSTEMS
- NORTHROP GRUMMAN
- HARRIS
- BOEING

©2012 DataDirect Networks. All Rights Reserved.

Massive Scalability Matters

DataDirect NETWORKS
INFORMATION IN MOTION

2008

Argonne NATIONAL LABORATORY

40,960 Nodes
8PB, 96GB/s

2011

Leading
Hedge Fund

2009

NATIONAL INSTITUTES OF HEALTH

6000 Node Cluster
600TB

2012

Leading
Analytics ISV
On Demand Service
for Drug Discovery

©2012 DataDirect Networks. All Rights Reserved.

DDN | Storage Fusion Architecture (SFA)

DataDirect NETWORKS
INFORMATION IN MOTION

Accelerating Big Data and Cloud, Optimizing TCO

Over 1 Million Lines of S/W Code – First Customer Shipped 2008
Designed Specifically for Big Data and Cloud Workloads

DirectMon™: Data Management Tools

*f*Scaler File System Family | *h*Scaler:Hadoop Appliance (tba)

In-Storage Processing™ Engine

Storage Fusion Architecture™
[Core Storage S/W Engine]

- Low-Latency Connect: FC, IB, Memory
- Interrupt-Free Storage Processing
- ReACT™ Adaptive Cache Technology
- SFX™ Flash Acceleration (tba)
- Quality of Service Engine
- Storage Fusion Fabric™

Parallel State-Machine Design
Maximum Performance, Lowest Latency

Virtualized Processing
Optimized Environment for Big Data Application Hosting

Robust Data Protection
Quality of Service and Performance Without Compromise

Flexible & Massively Scalable
Best-In-Class Scalability and Density

©2012 DataDirect Networks. All Rights Reserved.

New Enterprise Big Data Tools, Powered By Proven HPC Infrastructure

DataDirect NETWORKS
INFORMATION IN MOTION

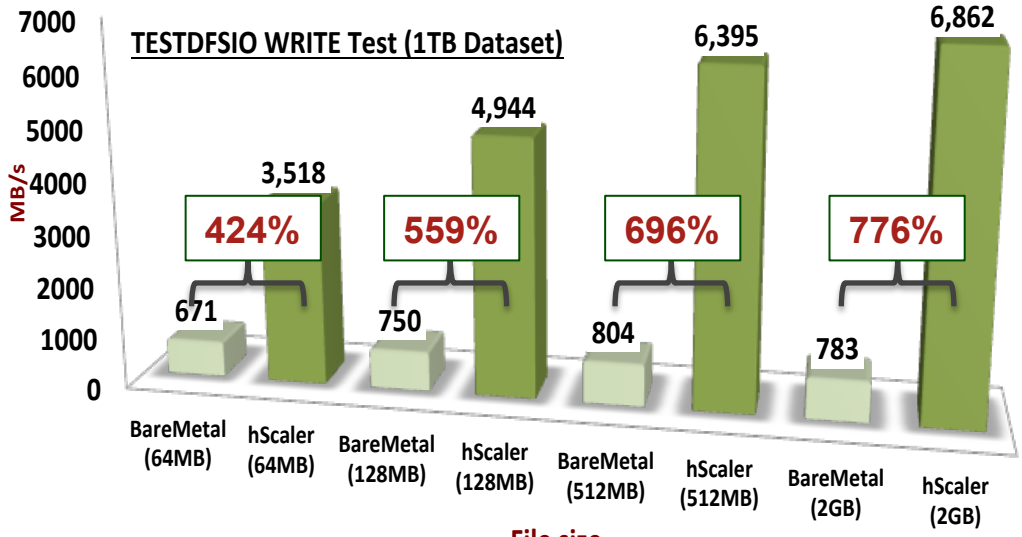

hScaler™

The DDN HPC Playbook, for Analytics:

- Full High-Throughput RDMA I/O
- Up to 700% Faster Wall-Clock
- Reduce Data Center Footprint By 60%

©2012 DataDirect Networks. All Rights Reserved.

hScaler Exemplifies A Broader Theme

HPC is Now Powering Enterprise Big Data!

DataDirect NETWORKS
Information in Motion

VERTICA
An HP Company

*1-Trillion Row Big Data Queries in **less than 20s**.*

SAS
THE POWER TO KNOW.

***Best Runtime Ever** for Drug Discovery, Warranty, Risk Analytics*

hadoop

*Enterprise Hadoop With **200%-700%** Performance Gain*

kx

*Up to **570% faster** FSI back-testing and risk management*

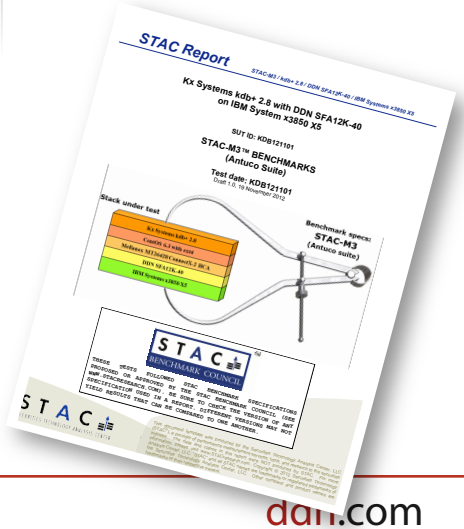

©2012 DataDirect Networks. All Rights Reserved.

The Power of Hybrid Storage, Today.

A Simple, Current Performance Case Study

DataDirect NETWORKS
Information in Motion

Without SFX

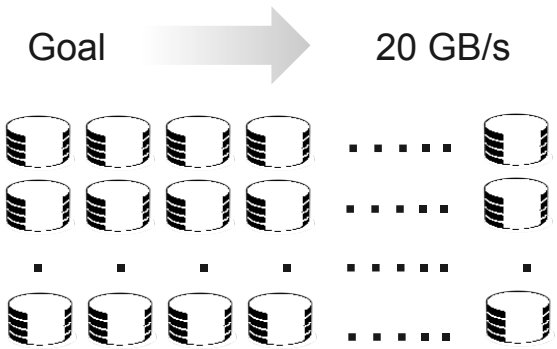

400 NL SAS Drives

With SFX

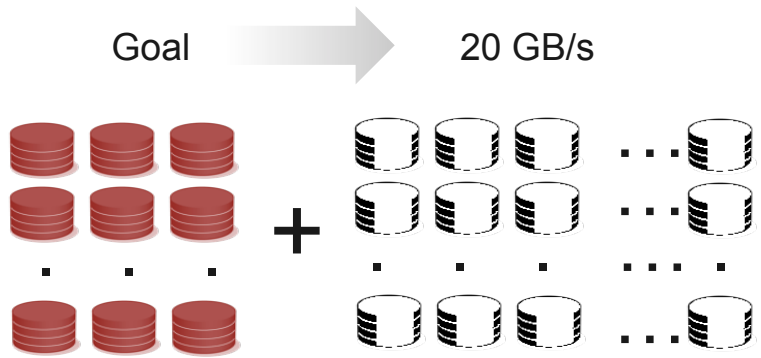

40 SSDs + 200 NL SAS Drives

	Mono	Hybrid	Gain
Drives	400 HDD	40 SSD; 200HDD	-
Power	4,400W	2,420W	45% Power Consumption Gain
Data Center	28U	16U	42% Reduction in Footprint
Cost (SRP)	$496K	$379K	25% Cost Advantage

As NVRAM Prices Decline & Concurrency Compounds, The Benefits of Hybrid Grow

©2012 DataDirect Networks. All Rights Reserved.

Cost/Latency Tradeoff

There is no single 'best' placement for BB NVM

DataDirect NETWORKS
Information in Motion

- Depending on workload and budget, any number of options exist for implementing a Burst Buffer layer that resides between CN and PFS
  - Interfaces may include DDR3/4, PCIe3, NVMe, SAS, SATA, etc.
- A robust burst buffer software stack must be adaptable to a wide variety of hardware implementations

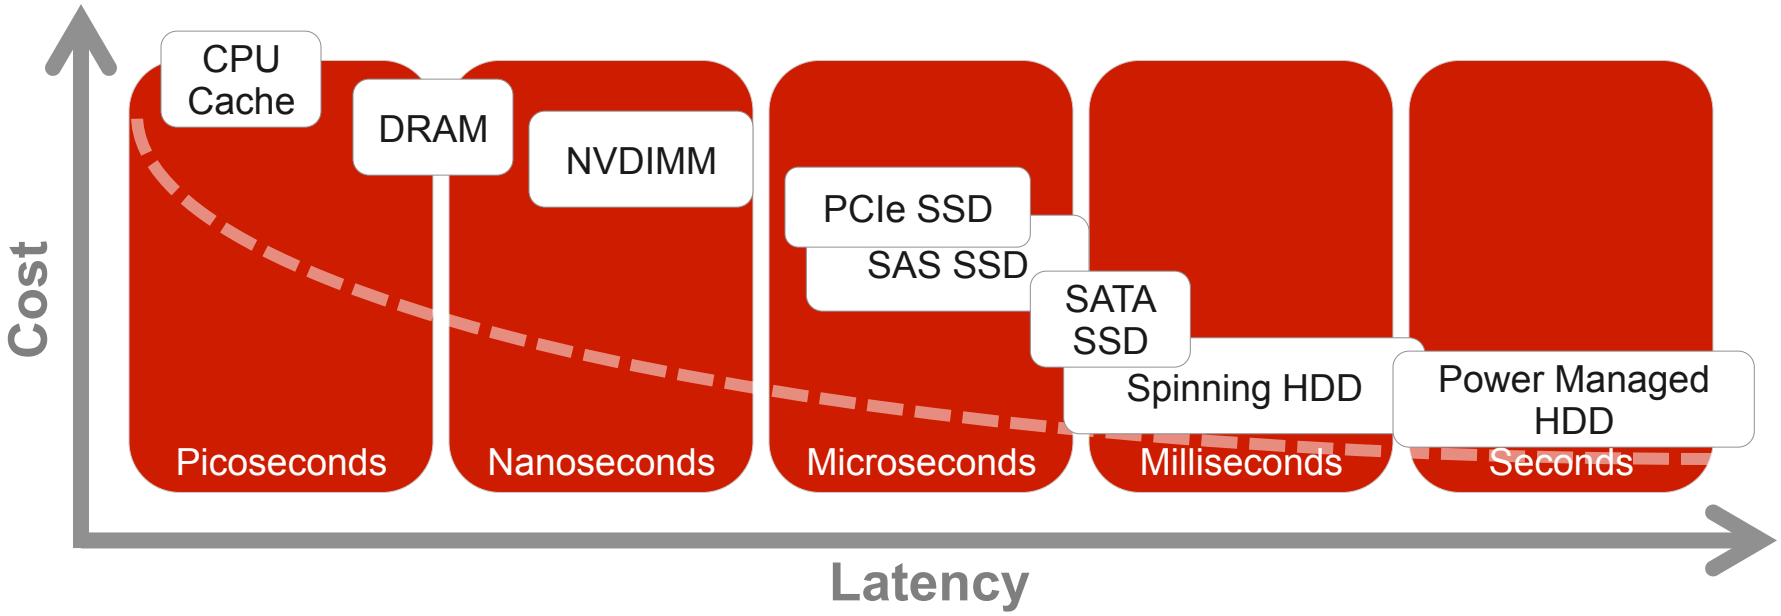

©2012 DataDirect Networks. All Rights Reserved.

We're Just Getting Started…

DataDirect NETWORKS
INFORMATION IN MOTION

	Buffer	Parallel File System	Archive / Cloud
Software Platform	Burst Buffer IRON MONKEY	Parallel File Storage lustre Storage Fusion Architecture Storage Fusion Xcelerator	Active Cloud-Enabled Archive WOS
Hardware Platform	NVM TBA	SFA: Hybrid Flash Arrays 	WOS Object Storage Appliances 

©2012 DataDirect Networks. All Rights Reserved.

DDN Global Collaboration Framework

DataDirect NETWORKS
INFORMATION IN MOTION

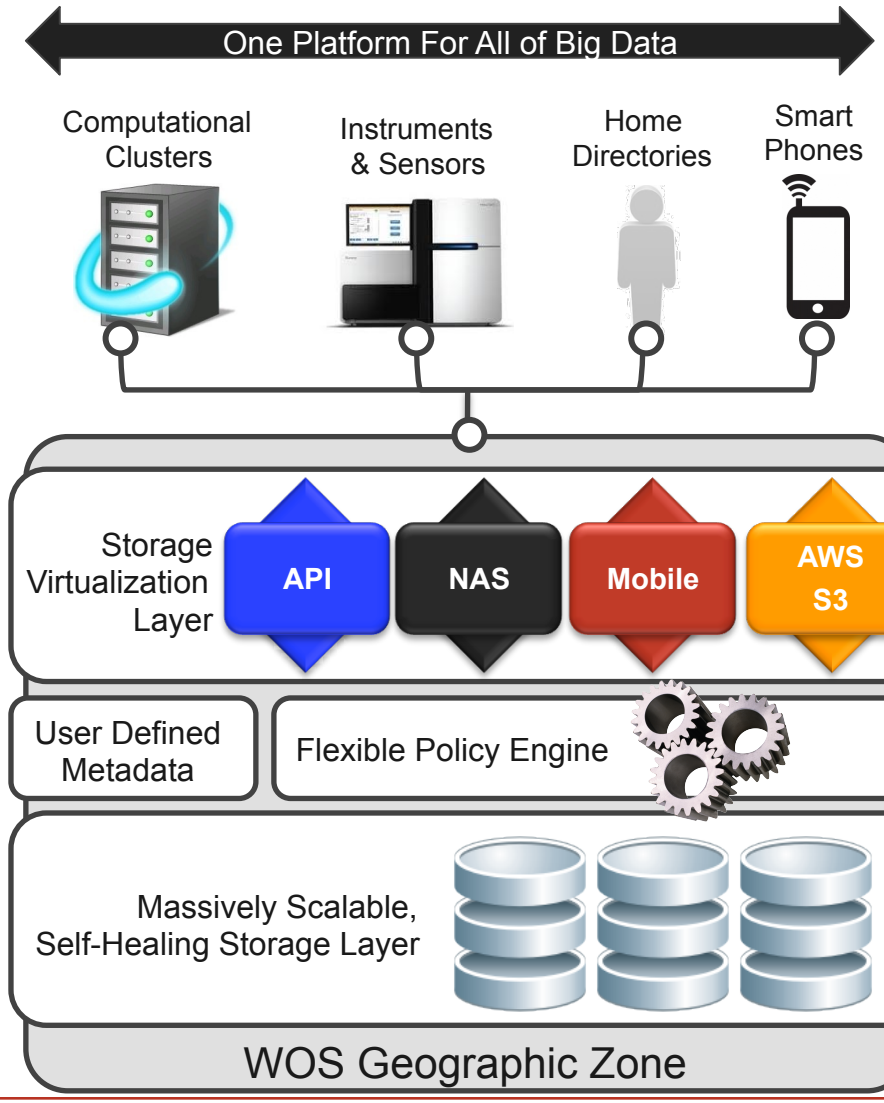

WOS for Data-Driven Collaboration

- Worldwide namespace access to massively scalable data volumes
- Secure cloud, behind your firewall
- Replication for DR or Collaboration to dozens of facilities worldwide
- Simple, self-healing infrastructure minimizes Big Data TCO
- World leading performance & latency

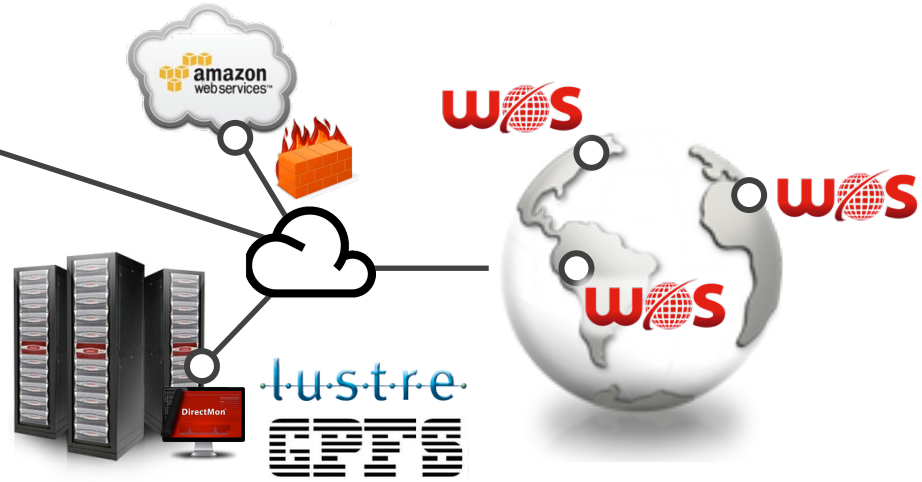

©2012 DataDirect Networks. All Rights Reserved.

WOS Enterprise Connectivity Options

DataDirect NETWORKS
INFORMATION IN MOTION

WOS Core

WOS API
Java/C++/Python

WOS Access

CIFS

NFS

WOS Cloud

Amazon S3

Desktop Sync & Web Client

WOS Share

Android Client

iOS Client

©2012 DataDirect Networks. All Rights Reserved.

End-to-End Architecture

DataDirect NETWORKS
INFORMATION IN MOTION

Buffer + FS + Archive + Cloud

A Fully Integrated Exascale I/O Platform To Minimize The Cost of Big Data Computing & Real-Time Analytics

Our opportunity resides in addressing the end-end efficiency and scalability challenge at 10^{18}…

… we're thinking **BIG**! Stay Tuned.

©2012 DataDirect Networks. All Rights Reserved.

DataDirect™
NETWORKS
INFORMATION IN MOTION™

Questions?

©2012 DataDirect Networks. All Rights Reserved.