



Flash Memory Summit

# Software-defined Storage and the Customer Path to NVMe-oF Parts 1 & 2

Ben Treiber

Sr. Director, Engineering Product Management



Flash Memory Summit

# 20 Years Helping Customers Meet Existing and Future Storage Needs

- >10,000 customers worldwide; enterprise proven
- 30,000+ installations across all popular server, storage & network manufacturers
- SDS deployments spanning several technology waves
  - Centralized, distributed, virtualized, Cloud, Hyperconverged
  - Interoperable across multiple generations and vendor brands
- Diverse & Mission Critical Apps; FC, NVMe, SSDs...
  - Tier-1 & database apps; Mfg / Healthcare / Fin / Gov't...
- 95% Customer satisfaction rating; 95% Renewal rate
- 10+ years of cash profitability + double digit growth



The real-time data company



Flash Memory Summit

# Why DataCore SDS? IT Challenges We Help Overcome



- Unacceptably slow applications due to I/O bottlenecks
- Interruptions from storage-related outages
- Old & New Co-existence – Evolutionary versus Disruptive Investment Protection plus Future-proofing
- Reduce cost & complexity of managing separate silos / islands of storage
- Doing more - Productivity and Performance

Santa Clara, CA  
August 2018



# Why Software-defined Storage?

The right software must be able to do a few things...

**1** Enable different storage devices to communicate with one another

**2** Separate advances in software from advances in hardware

**3** Pool all storage capacity and provide centralized management

**4** Make hardware maintenance, data migrations, and refreshes non-disruptive and easy

**Plus Enhance Business Continuity, Performance and Automation**





## Why SDS? – Customers report:

IMPROVE  
PERFORMANCE



79% improved  
performance by  
**3X or more**

REDUCE  
DOWNTIME



81% reduced storage-  
related downtime by  
**50% or more**

INCREASE  
UTILIZATION



82% reduced storage-  
related spending by  
**25% or more**

SAVE  
MONEY



100% saw, in the  
first year,  
**Positive ROI**

Research by  **TechValidate** from Survey of over 1000 DataCore customers around the globe



# Customer Feedback

## CASE STUDY: WESTERN DIGITAL AND DATACORE COMBINE FOR A PERFORMANCE DISRUPTIVE SOFTWARE DEFINED STORAGE SOLUTION

- Current joint customer scenario deployed at a major service provider:
  - Existing Storage based on 2x EMC Vplex systems each equipped with 90 TB of traditional HDD storage;
  - Backed up by a VNX system with a SSD Fast Cache.
  - The Infrastructure was ageing, support was expiring, and needed more performance and a better TCO over 5 years.
  - Budgets were tight.
- **A Software Defined Storage solution supporting adaptive auto-tiering was needed, resulting in an all flash system delivering 3X the performance against competitive AFAs but made up of cost-effective Western Digital HGST NVMe SSDs and 2U24 JBOFs.**
- NVMe SSDs and JBOF storage + DataCore SANsymphony SDS Auto-Tiering (Economic Optimization) = **Flexible and Affordable All Flash Solution**



# Customer Feedback



## **CASE STUDY: DATACORE DELIVERS TO MACSTADIUM AND RIDE-SHARING CUSTOMER 3-5 TIMES THE VALUE OF TRADITIONAL ALL-FLASH ARRAY STORAGE**

A major customer, a leading ride-sharing company relies on MacStadium to create short-lived virtual machines for Mac iOS build environments. **Virtual machines spin up and down continuously** within a matter of minutes—all day, every day. As **the number of VMs began to grow**, the company's workload began to **tax the I/O performance of MacStadium's existing storage solution.**

**MacStadium had two options: Adding more all flash arrays or introducing software-defined storage.**

MacStadium chose the latter and **deployed DataCore SANsymphony SDS with Parallel I/O.**

**“Despite our best efforts to break DataCore, it has continued to work without fail, proving its reliability. And now that it's been running in our customer's environment for over one year, it's clear that DataCore provides three to five times the performance of traditional all-flash arrays at the same price point.”**

Santa Clara, CA  
August 2018

-Jason Michaud, founder and president of MacStadium



# Customer Feedback



## **CASE STUDY: DATACORE DELIVERS TO MACSTADIUM AND RIDE-SHARING CUSTOMER 3-5 TIMES THE VALUE OF TRADITIONAL ALL-FLASH ARRAY STORAGE**

**“DataCore SDS delivers far more flexibility than other solutions,” he said. “And because it includes a unified storage services management interface and synchronous mirroring, we can perform maintenance—planned and unplanned—whenever we choose.”**

MacStadium has been so pleased with DataCore that it expanded use of the software to an additional eight clusters.

**“Many of our customers have demanding workloads, but DataCore is tunable, giving us a reliable, set and forget solution. DataCore is our sledgehammer solution because it addresses everything our customers require.”**

-Jason Michaud, founder and president of MacStadium



# Common Customer Challenges With NVMe Storage and NVMe-oF

## INTEGRATION AND MANAGEMENT

- How do I integrate this new technology with my existing environment?
- What new processes do I need to manage this investment?

## AVAILABILITY OF SOFTWARE FEATURES

- How do I get the functionality I need like high-availability, failover and data protection?
- What tools will I have to handle data migrations, QoS, tiering and thin provisioning to ensure that I'm making the most of my investment?

## SHARING FLASH STORAGE BETWEEN APPLICATIONS

- Do I really need to buy dedicated NVMe cards for every application?
- How can I share and leverage my existing storage investments along with new NVMe storage?



## NVMe-oF + SDS Use Cases?

- Initial Feedback from DataCore's customer base on use cases:
  - **Co-existence with existing storage** - Use NVMe over Fabrics to co-exist with installed Flash devices and storage subsystems that use traditional protocols internally to handle I/O. Customers gain the benefit of using a proven data services software stack and lower latency over the wire, while taking advantage of existing storage subsystem technology.
  - **New system deployments** - Pooling and sharing NVMe devices (JBOFs) for cost savings versus AFAs - Customers using a storage system comprised of many NVMe devices typically for ERP / Database applications, using NVMe over Fabrics with a Fibre Channel interface, making a complete end-to-end high performance, low latency NVMe storage solution.



Flash Memory Summit

# Recent Collaboration Announcement

## DataCore and Broadcom to Accelerate Adoption of NVMe Storage

- **“Unlike other NVMe solutions, which have deployment limitations or require new network fabrics and ‘rip and replace’ of current investments, ours is essentially a plug-and-play solution that will enable companies to immediately begin realizing the significant benefits of NVMe with a path to NVMe-oF.”**


-Steven Hunt, product management leader at DataCore.

- **“DataCore and Broadcom provide customers with the fastest Fibre Channel connectivity combining low-latency and highest IOPS performance with NVMe over Fibre Channel support to meet today’s business-critical workloads, as well as enabling next-gen applications to run at scale on existing SAN infrastructure. **Broadcom can run traditional SCSI FCP concurrently with the NVMe over Fibre Channel command set on the same Fibre Channel network, making the transition to the all NVMe datacenter seamless.**”**

- Jeff Hoogenboom, general manager of the Emulex Connectivity division at Broadcom.

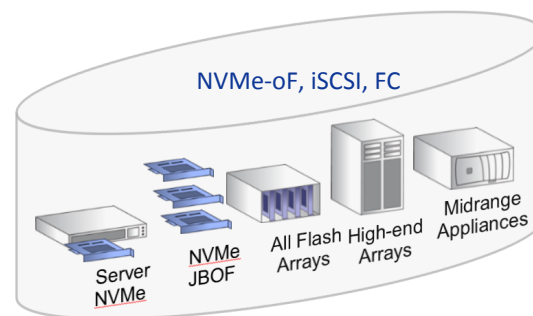
## 'End-to-end NVMe-oF' Word of Caution

Many vendors are touting end-to-end NVMe-oF but they really mean all roads lead to their vendor supplied 'No co-existence' models. What about existing storage investment protection?



"Storage Island"

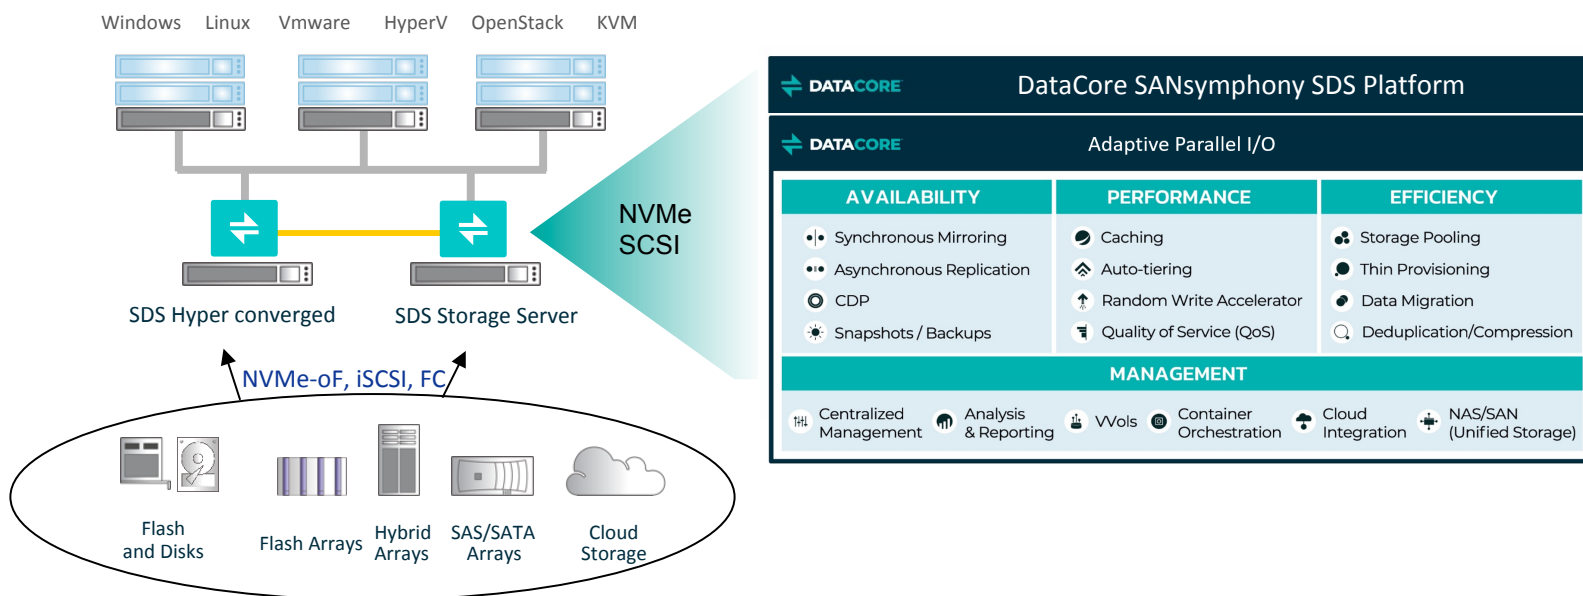
Flexible SDS solutions offer end-to-end NVMe-oF supporting any storage and the ability to leverage not only a mix and match of AFAs but any NVMe device / JBOF or other storage devices.



"Choice"







# DataCore SDS Overview



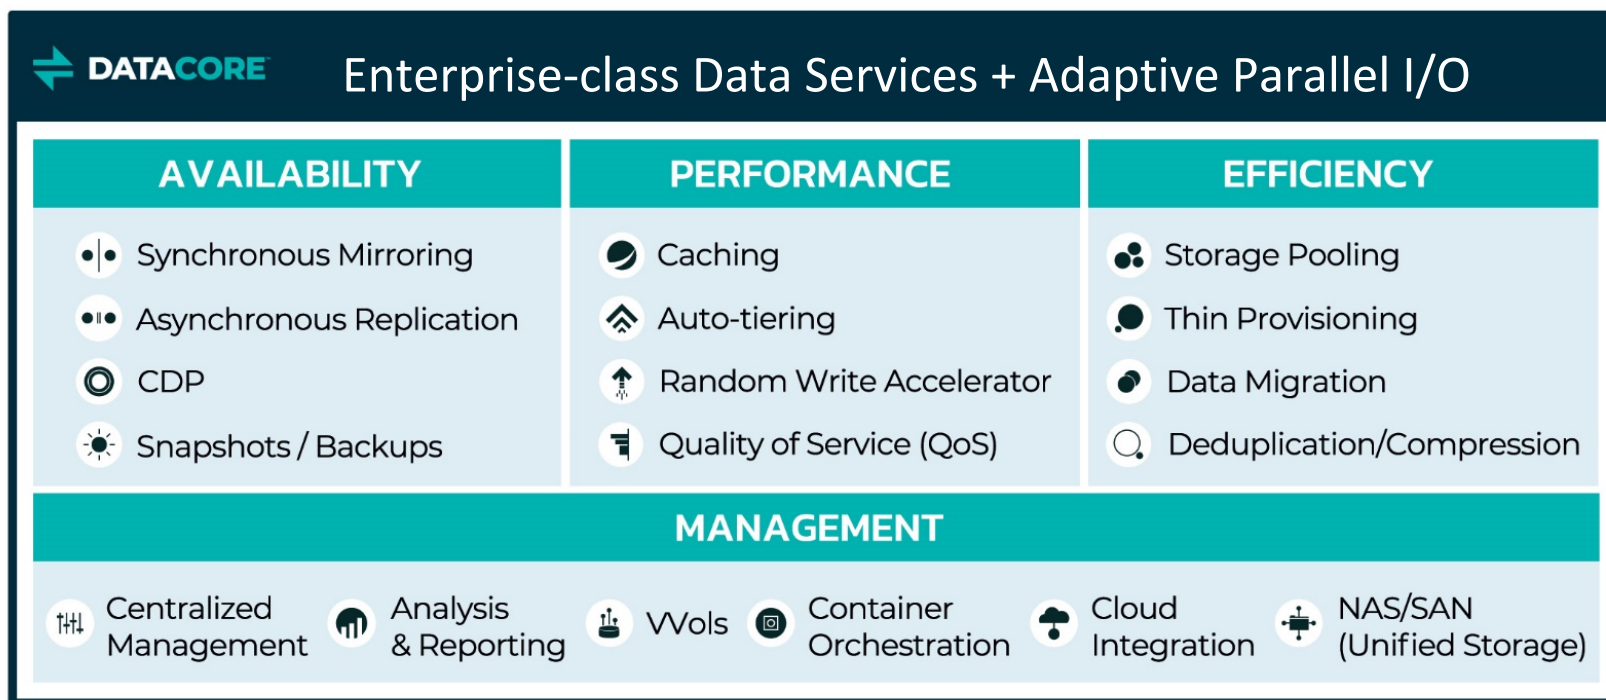
- Any standard x86 servers/VMs in the data path; Any Hypervisor on Linux / Windows
- Any storage; Pool & Share NVMe, SSDs, AFAs, HDDs...Route Choice NVMe or SCSI

# Flexible SDS Deployment Models




Traditional Storage Virtualization	Converged (Server SAN)	Hyper-Converged	Hybrid-Converged
Integrate, manage, and enhance existing storage	Leverage internal storage, reduce complexity and maintain compute segregation	Consolidate all functions for smallest footprint and highest performance	Consolidate all functions for smallest footprint and highest performance while serving storage externally

# NVMe-oF Needs Proven Software to Accelerate Customer Adoption



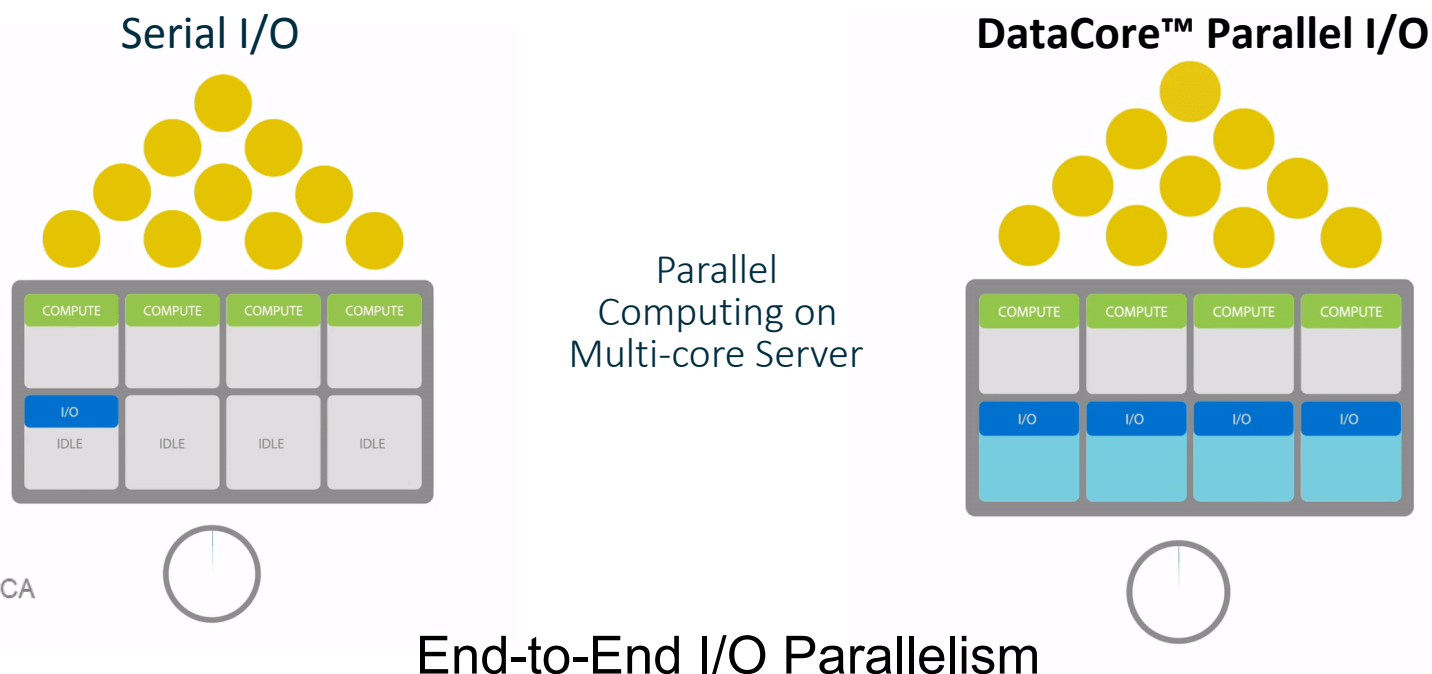
# Continuous Availability & Data Protection No Downtime Refresh / Data Migration

Move contents from one device to another without impacting applications



- Zero Touch Failover / Failback
- HA sync mirrors span metro-wide
- Eliminates downtime for migrations, refreshes & upgrades
- No disruption to Apps
- Virtual disks remain undisturbed
- Transparent Data Migration
  - Transfers device/array contents to new equipment in background

# DataCore Parallel I/O Removes serial I/O bottleneck



Santa Clara, CA  
August 2018

## End-to-End I/O Parallelism

NVMe takes advantage of parallel, low latency data paths to underlying media  
DataCore Parallel I/O takes advantage of multi-cores to remove serial I/O bottlenecks prior to controller / data paths.




# DataCore Parallel I/O SPC-1™ OLTP Workload Benchmark

Based on industry-standard SPC-1™  
see Full Disclosure Reports in notes

#	SPC-1 IOPS™	Tested Product	Avg Response (ms)	Price	Price-Performance
1	5,120,098	DataCore Parallel Server 2-node	0.28	\$ 506,525	\$0.10
2	3,010,007	Huawei OceanStor 18800V3	0.92	\$ 2,370,763	\$0.79
3	2,004,941	Hitachi VSP G1000 / HP XP7	0.96	\$ 2,003,803	\$0.98
4	1,239,898	Kaminario K2 (K2F00000700)	2.95	\$ 997,348	\$0.80
5	1,201,961	DataCore SANsymphony HA-FC	0.22	\$ 115,142	\$0.10
6	1,005,893	Huawei OceanStor 18800	5.17	\$ 2,794,971	\$2.78
7	780,081	IBM Power 780 server with SSDs	18.9	\$ 3,557,709	\$4.56
8	685,281	NetApp FAS8080 EX All-Flash	1.23	\$ 1,897,999	\$2.77
9	650,987	Huawei OceanStor 6800 V3	3.36	\$ 1,488,036	\$2.29
10	605,016	NEC Storage M710F	1.71	\$ 492,726	\$0.81

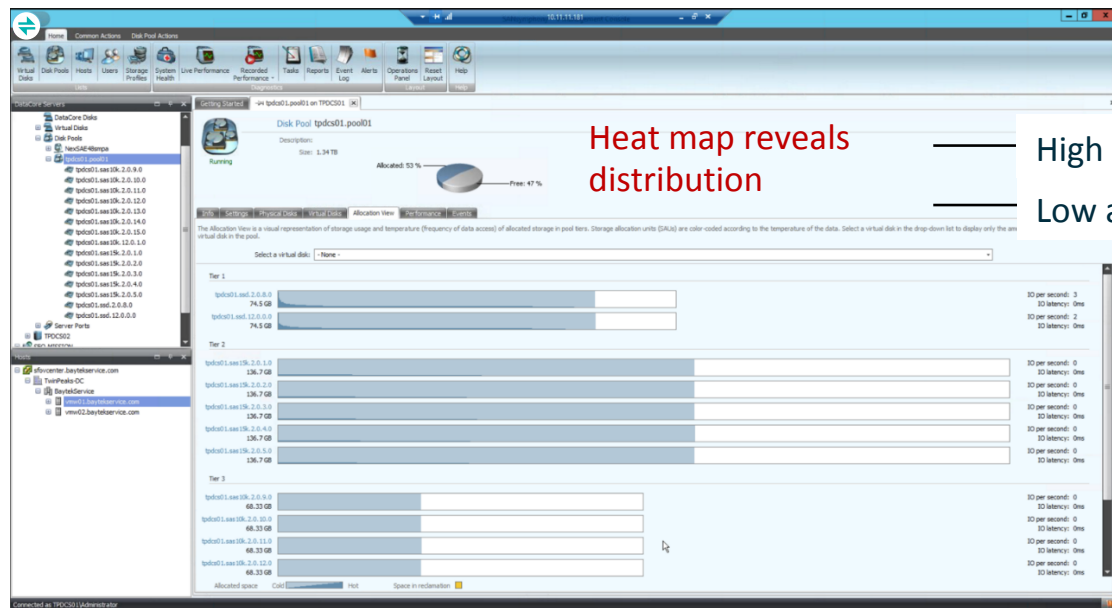
# Adaptive Auto-tiering Automatic Resource Optimization



Evaluates data at the block level so data from a single application  
can reside on multiple storage tiers

# Adaptive Tiering - Automatic Tuning


Blocks spread across storage resources within defined tiers



Heat map reveals distribution

High activity

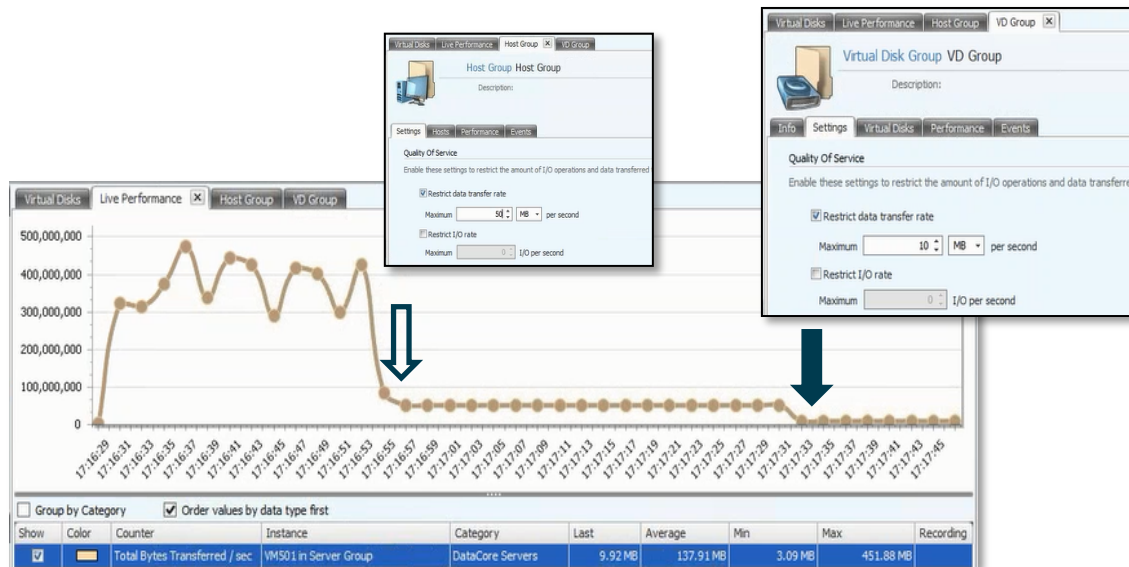
Low activity





# QoS Controls

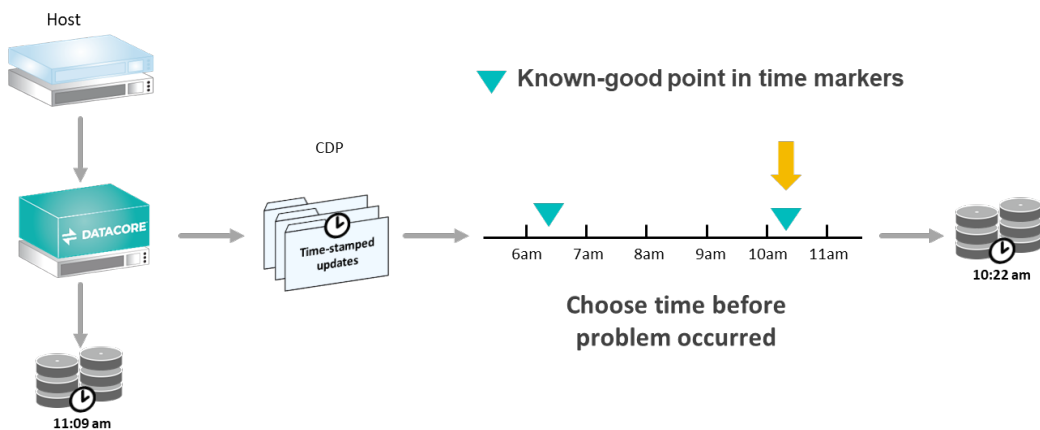
## Regulate IOPS and Throughput



- Enables critical apps to run faster
- Set limits on I/O traffic generated by less important workloads
- Saves money otherwise spent on dedicated equipment

# Continuous Data Protection (CDP)

Return to any point-in-time without taking explicit backups




- Restore arbitrary point-in-time
- Logs and timestamps all I/Os to the selected virtual disks
- No need to quiesce or interrupt applications
- No host agents required
- Easy to enable and revert from
- Create known-good restore points



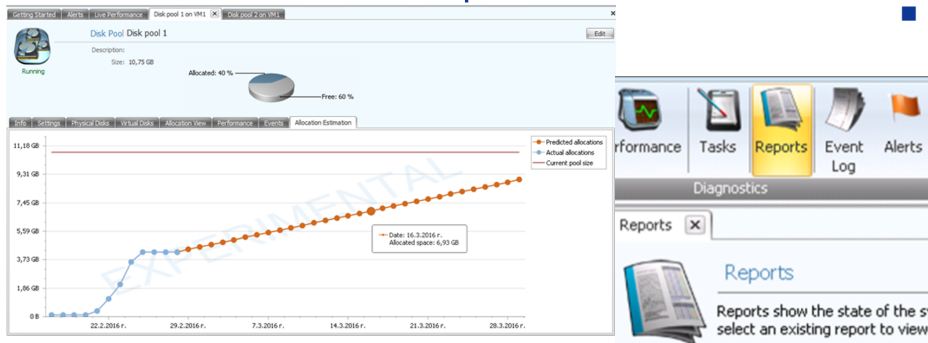
Flash Memory Summit

# Analysis & Reporting

Capacity planning, performance tuning...



Predict Future Consumption



- Configurable views of system behavior and performance
- Capacity and Performance Trending
- Drill down into I/O characteristics
- Export resource allocation to billing & chargeback tools

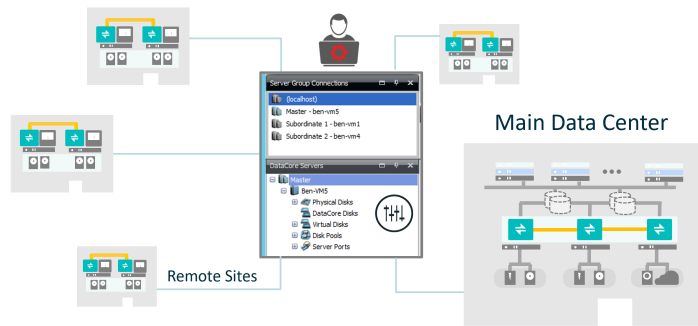
Santa Clara, CA  
August 2018



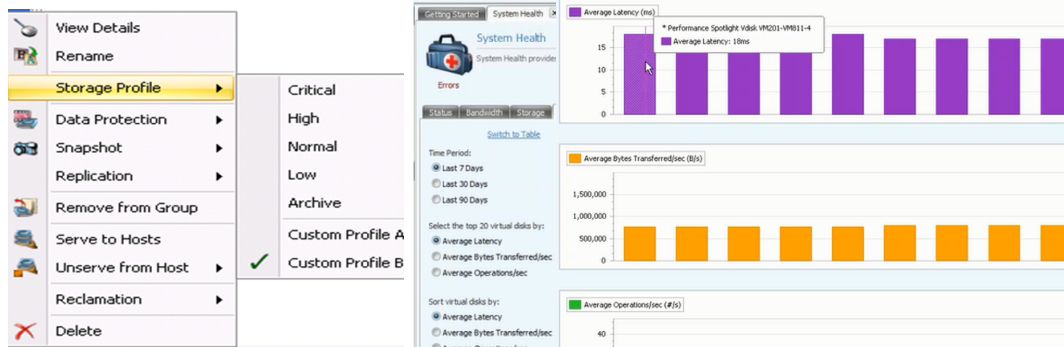
Flash Memory Summit

# Centralized Storage Management

Control & monitor diverse resources



Where is Attention Needed?



- Intuitive to set up and operate
- Automates repetitive tasks
- Templates, Profiles, Self-guided wizards for key workflows
- Comprehensive diagnostics
- Role-based, administrative permissions & audit trail
- Integrated with popular hypervisors / containers
  - Vmware / Hyper-V / KVM
  - Docker/ Kubernetes



## Demonstration Story Board

- DataCore can demo **SDS data services** over NVMe-oF to hosts running either:
  - LINUX
  - WINDOWS
- Demo is focused on Windows initiator and our **target driver allowing any storage backend** since it is a major gap in the marketplace



## Demonstration Story Board – Scenario

- Create an NVMe Namespace (vDisk) for presentation to an NVMe-oF capable host
- Will present an NVMe-oF Namespace (vDisk) to a Host
- With DataCore SDS, this is simple!

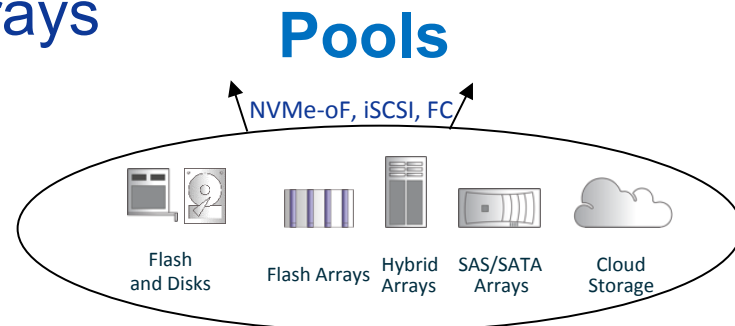


## Creating and preparing a NVMe Namespace (vDisk)

- SANsymphony SDS manages any physical storage device capable of presenting a storage device to Windows Logical Disk Manager (LDM)
- Any storage device Windows sees, DataCore SDS can share and manage

# Creating and preparing a NVMe Namespace (vDisk) – cont'd

- Managed devices can be internal or external:
  - NVMe
  - PCIe
  - SAS
  - SATA
  - FC
  - IDE
  - All Flash Arrays (AFA)
  - Just a bunch of Flash (JBOF)
  - Hybrid Arrays
  - JBOD
  - SSD
  - HDD

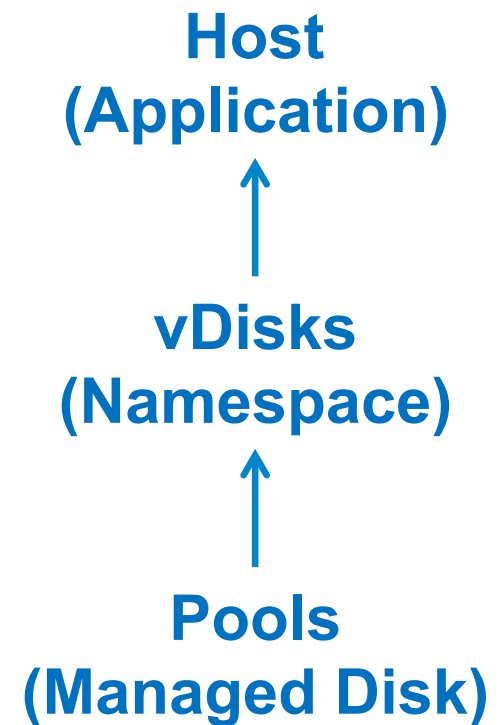






## Creating and preparing a NVMe Namespace (vDisk) – cont'd

- Pools of storage devices are defined
  - These are managed by DataCore SDS software
- vDisks are
  - Logical devices created in the software with attributes (such as NVMe-oF, HA)
  - Attributes and data services per vDisk (Adaptive tiering, Data Migration, QoS, Load balancing...)
  - Presented to Hosts via FC targets that implement the NVMe-oF protocol

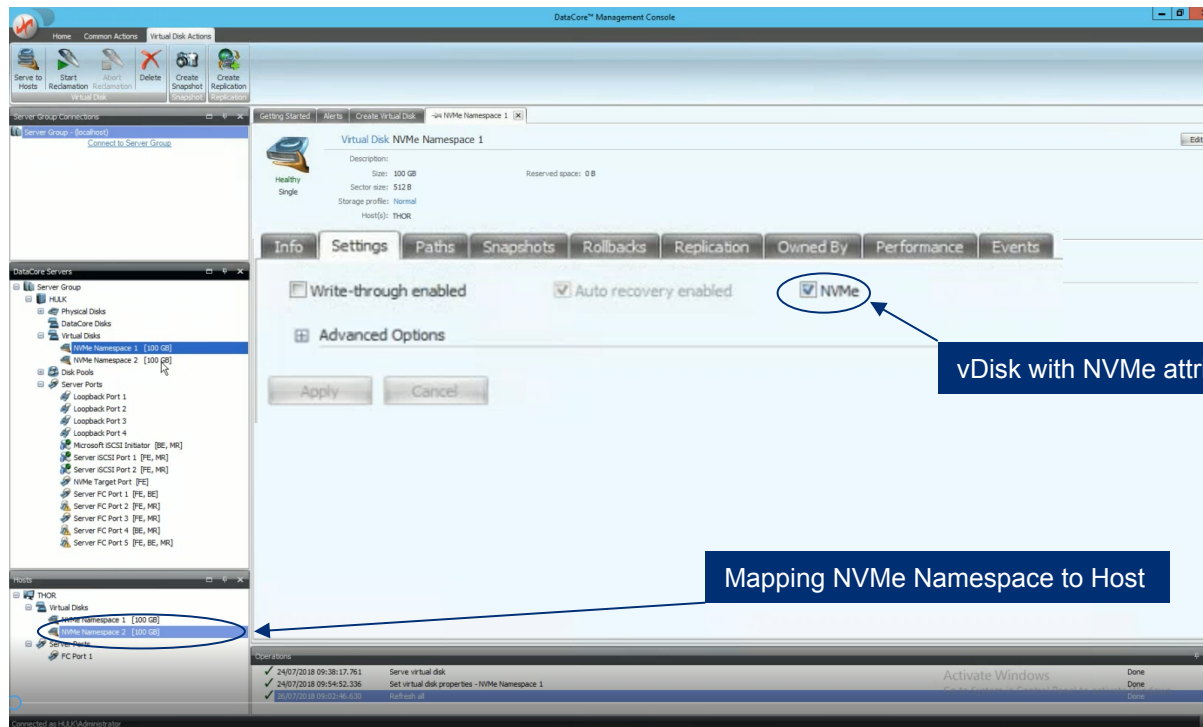




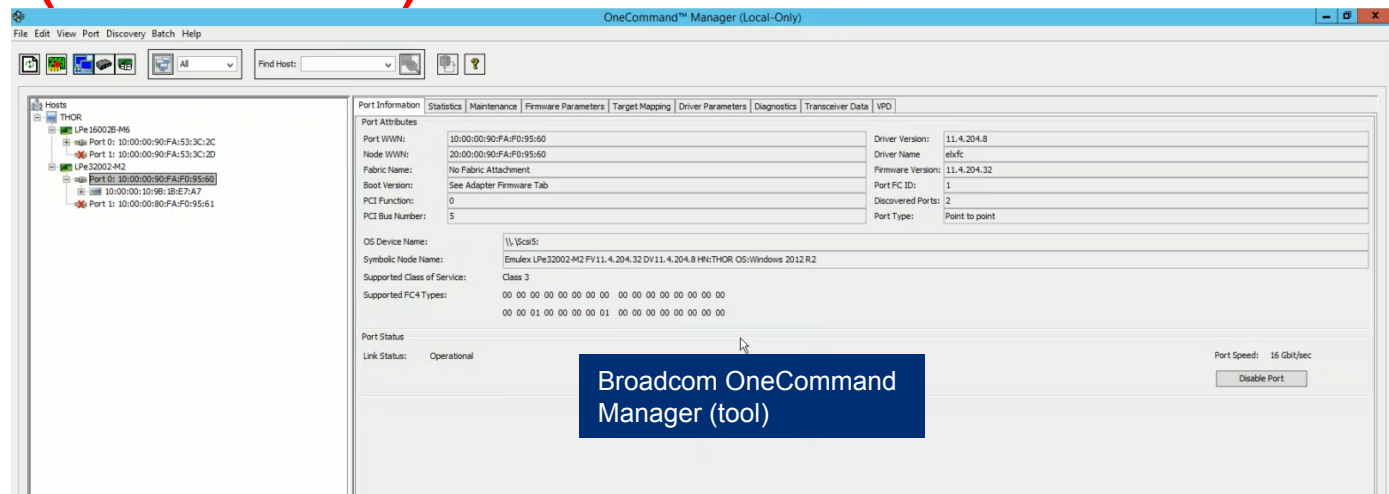
# Creating and preparing a NVMe Namespace (vDisk) – cont'd

- Our Demo Scenario
  - Windows Server
    - LPe32002 FC Host Bus Adapter (FC Initiator)
  - Connected to DataCore SDS
    - LPe32002 FC Host Bus Adapter (FC Target)

# Creating and preparing a NVMe Namespace (vDisk) – cont'd

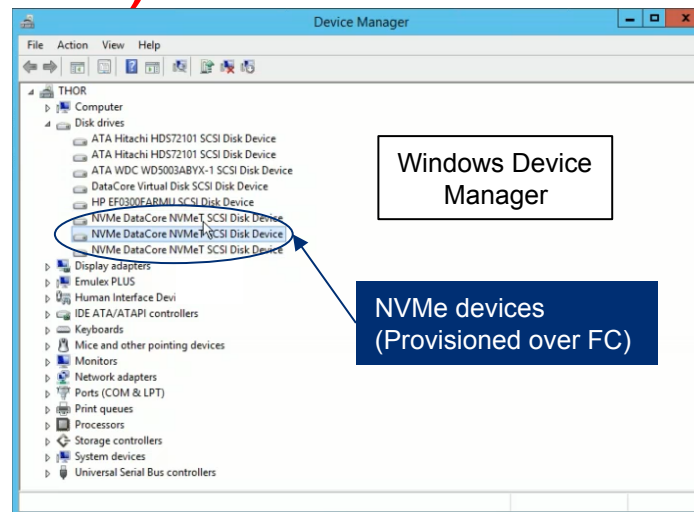


# Discovering provisioned devices (host side)



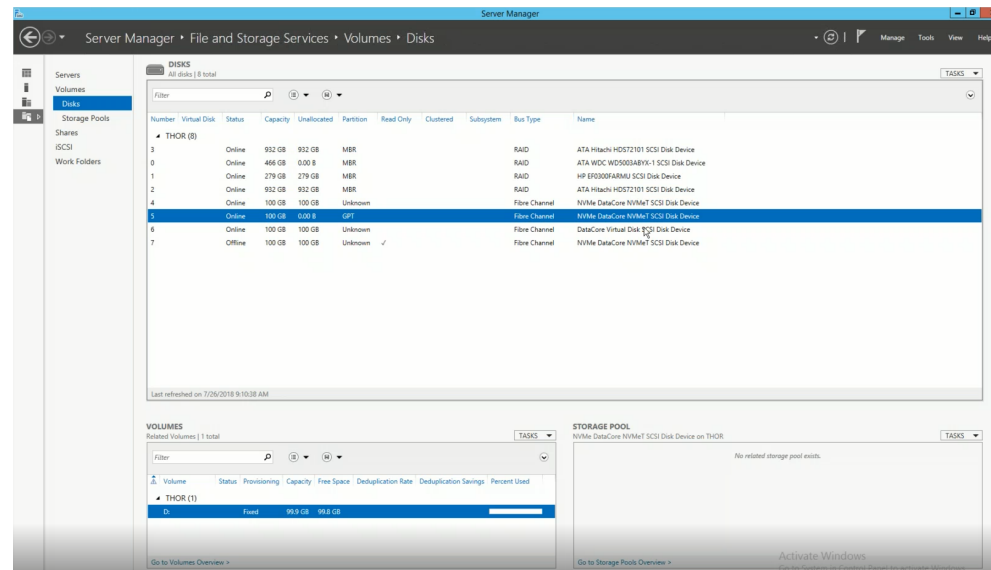
- Collaboration with Broadcom
- Used OneCommand tool to discover NVMe-oF target devices for Windows

## Discovering provisioned devices (host side)



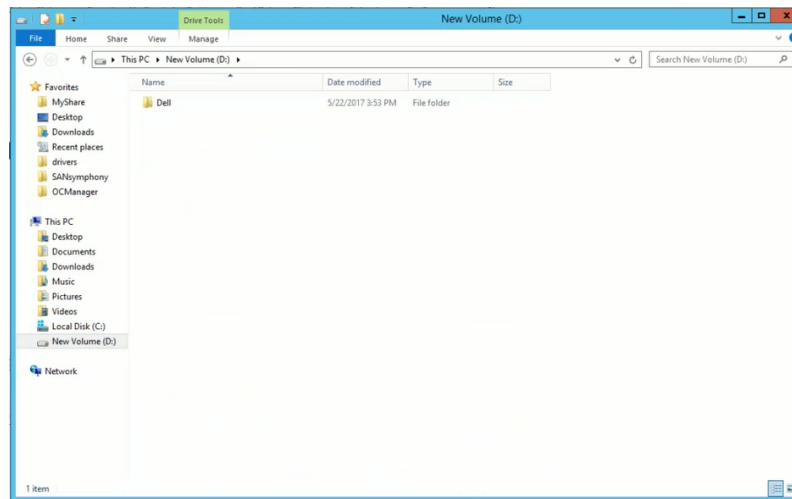
- Here, Windows Device Manager displays the discovered NVMe devices which are treated as Disks, like any other but faster...simply, NVMe devices provisioned by SDS

# Partitioning / Formatting provisioned devices (host side)



- Partition / Format the volume in the usual way with Common tools - using Microsoft Server Manager

## Using provisioned devices (host side)



- The NVMe-oF disk is well behaved (use it like any disk with a filesystem) – only more responsive and faster!



# Performance Over the Wire – with services! – Wire Trace

~14 Microseconds (NVMe-oF)

Boo...	Side A	Side B	Age(us)	Pend...	LBA	Expe...	Tag	LUN	Source	Destination	mm:ss.ms_us_ns_ps	Delta
	Read	Good	15.842_9	1	3DB0	1000	03F2				00:00.667_202_224_0	
	Read	Good	16.641_3	2	3DB8	1000	03E3				00:00.667_216_597_2	
	Read	Good	16.857_4	3	3DC0	1000	03EE				00:00.667_217_839_2	
											00:00.667_219_081_4	
	Read	Good	16.167_4	4	3DC8	1000	03EC				00:00.667_223_524_0	4
	Read	Good	16.454_5	4	3DD0	1000	03E5				00:00.667_229_039_0	5
	Read	Good	15.982_1	4	3DD8	1000	03E1				00:00.667_232_395_0	3
	Read	Good	16.505_2	4	3DE0	1000	03F2				00:00.667_238_758_0	6
	Read	Good	17.005_9	4	3DE8	1000	03E3				00:00.667_244_302_0	5
	Read	Good	16.070_1	4	3DF0	1000	03EE				00:00.667_248_222_0	3
	Read	Good	16.662_0	4	3E00	1000	03E5				00:00.667_258_346_0	10
	Read	Good	21.828_2	4	3DF8	1000	03EC				00:00.667_258_794_0	0
	Read	Good	15.617_0	4	3E08	1000	03E1				00:00.667_262_125_0	3
	Read	Good	16.143_9	4	3E10	1000	03F2				00:00.667_267_683_0	5
	Read	Good	15.795_7	5	3E18	1000	03E3				00:00.667_272_290_0	4
	Read	Good	16.616_8	4	3E20	1000	03EE				00:00.667_278_087_0	6

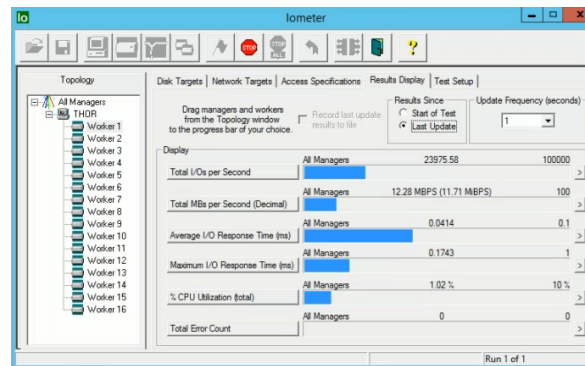
Analyzer trace – target read command, NVMe-oF

Note: Tech Preview, current focus has been functionality and stability versus performance optimization. Next step performance optimization...



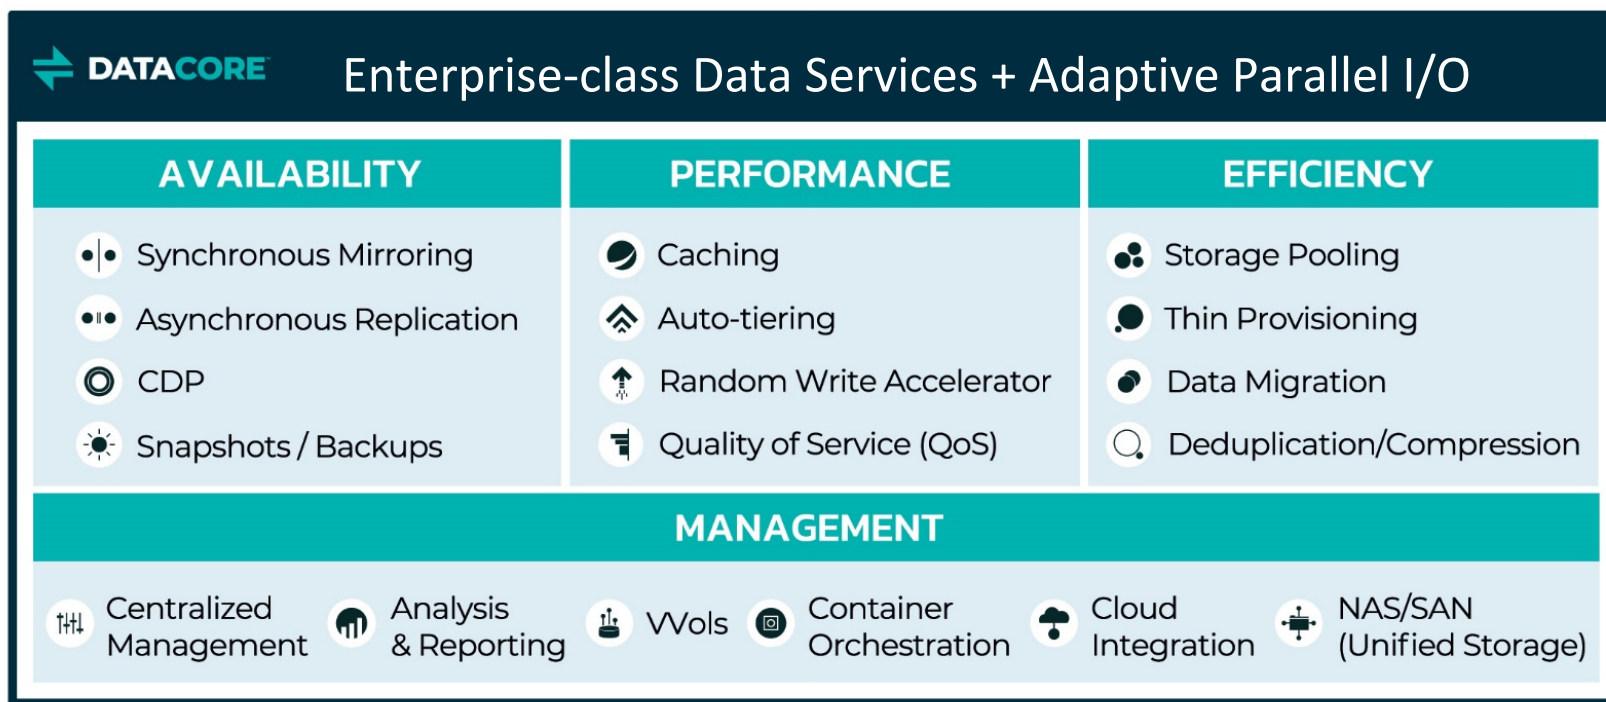
# Performance Over the Wire – with services! – User Level IOmeter Run

Windows  
~41 Microseconds  
(NVMe-oF)



Note: Tech Preview, current focus has been functionality and stability versus performance optimization. Next step performance optimization...

# NVMe-oF Needs SDS to Accelerate Customer Adoption





Flash Memory Summit

## Key Take-aways

- Flexible SDS solutions **Improve** Performance and Utilization; **Reduce** Down-time, Cost and Management Complexity
- NVMe-oF is the next **Evolution** for Performance and Lowering Latency
  - DataCore **Parallel IO** unlocks NVMe-oF and FC-NVMe's full performance
- NVMe-oF needs **Data Services** to ease Enterprise deployments
  - Same tools, QoS, auto-tiering, CDP, replication, load balancing, etc.
- An **Open SDS** platform is the **Bridge** that **Unifies** and **Abstracts** legacy (Investment Protection) and new storage; NOT create NVMe islands

**NVMe-oF + DataCore SDS Provides the Best of Both Worlds:  
Lowest Latency & Highest Performance + Enterprise Class Data Services**



Questions?



Thank You!