

Building Self-Healing Mass Storage Arrays for Large Cluster Systems

NSC08, Linköping, 14. October 2008

Toine Beckers
tbeckers@datadirectnet.com



Agenda



- **Company Overview**
- **Balanced I/O Systems**
- **MTBF and Availability**
- **Recovery Operations**
- **Parallel Storage Systems**
- **Conclusion**

DataDirect Leadership



- Established 1988
 - Technology Company (ASICs, FPGA, Firmware, Software)
 - System Experts (Optimization, Clusters, Interconnects, Protocols, File Systems, Streaming, Video)
 - S2A Introduced June 2000 (Developed from 1997 to 2000, Shipping 8th Gen)
- Focused
 - High Throughput, High Scalability
 - HPC and Media & Entertainment
- > 4,000 Systems Shipped
- 7th Generation S2A9550 in Q4'05
- 8th Generation S2A9900 in Q2'08
- Recent Gartner Dataquest report stated: DDN is 5th largest independent storage provider in terms of Market Share, and 3rd largest independent storage provider in volume.
- #2 Fastest Computer in the World
 - DDN Powers IBM's BG/L @ LLNL
 - S2A Delivers 320 TFlops w/ 1PB of SATA
- 6 from Top 10 and 36 of the HPC Top50 Sites
 - DDN Powers Clusters from IBM, Dell, HP, Cray, SGI, Bull, others...
- #1 Tapeless Newsroom in the World
 - DDN Powers CNN
- > 400 Postproduction and Broadcast facilities
 - DDN Powers Systems from Sony, SGI, Autodesk/Discreet, Pinnacle, Thomson, ...

Rank	Site	Manufacturer	Computer	File System
2	DOE/NNSA/LLNL	IBM	eServer Blue Gene Solution	Lustre
3	Argonne National Laboratory	IBM	Blue Gene/P Solution	pVFS + GPFS
5	Oak Ridge National Laboratory	Cray Inc.	Cray XT4 QuadCore 2.1 GHz	Lustre
6	Forschungszentrum Juelich (FZJ)	IBM	Blue Gene/P Solution	GPFS
9	IDRIS	IBM	Blue Gene/P Solution	GPFS
10	Total Exploration Production	SGI	SGI Altix ICE 8200EX, Xeon quad core 3.0 GHz	Lustre

Extreme Storage

Where to Position ?

DataDirect[™]
NETWORKS



PERFORMANCE

Throughput-Intensive Applications

>1.5GB/s Sustained Read & Write to 250GB/s

- Real-Time Streaming/Ingest
- Collaborative Data Access
- Bandwidth / Clustered Bandwidth

SCALABILITY

Scalable Capacity Applications

>100TB to Dozens of Petabytes

- Regulatory Compliance
- Backups, Archival & Virtual Tape
- Storage Consolidation, Content Stores



Sample HPC Partners & Customers



Over 100 GB/s of Extreme Storage at LLNL



Lawrence Livermore National Laboratories needed to maintain data reliability while deploying SATA technology and **Scale Over 100GB per Second with Consistency.**

Only S2A could deliver LLNL's **Extreme Storage** performance requirements.



EXTREME STORAGE

Number One Oil Producer on Extreme Storage

DataDirect[™]
NETWORKS



When you're the World's #1 Oil Producer,
You Need Extreme Storage. S2A Delivers.

أرامكو السعودية
Saudi Aramco



Saudi Aramco supplies **10% of all the oil consumed in the world**. They needed highly advanced storage technology to meet their **Extreme Seismic Exploration** needs.

With over 48 S2A systems from IBM, Saudi Aramco is equipped to accelerate time to oil.

Agenda



- Company Overview
- **Balanced I/O Systems**
- MTBF and Availability
- Recovery Operations
- Parallel Storage Systems
- Conclusion

Balanced I/O in Clusters

DataDirect[™]
NETWORKS



- **Balance is essential for minimizing I/O in cluster computing**
- **DOE has generally used 1GB/s/TF**
- **Current systems range to 300GB/s**
- **Petaflop scale clusters assume 1TB/s**
 - At 100mB/s/drive one system would contain 10,000 drives at 180kW for SAS

Disk Drive Progress

DataDirect[™]
NETWORKS



Cheetah 1 FC

- Dual ported at 100MB/s
- 1GB capacity
- Sustained reads at 5MB/s
- 6.5mS full stroke seek
- Block reassign in ~1.5s

Cheetah 7 FC

- Dual ported at 200MB/s
- 300GB capacity
- Sustained reads at 50+MB/s
- 6.5mS full stroke seek
- Block reassign in ~2.5s

Challenge: How to achieve dramatic performance increases with no change in disk random performance

Solution: High Performance
Silicon Based Storage Controller

- **Parallel access for hosts and parallel access to a large number of disk drives**
 - True performance aggregation and scalability
 - Reliability from a parallel pool and QOS
- Drive error recovery in real time and True State Machine Control

Challenges to I/O Balance

- **Current SAS drives still have full stroke seeks at 6.5mS**
- **Current SATA drives still have full stroke seeks at 22mS**
- **Recovery mechanisms are slowed by increased data density**
 - SAS drives take 75 steps
 - SATA drives take 300 steps
- **Block reassign times can range to tens of seconds including recovery**
 - SAS drives at 1.5s/LBA
 - SATA drives at 6s/LBA

Agenda



- Balanced I/O Systems
- **MTBF and Availability**
- Recovery Operations
- Parallel Storage Systems
- Conclusion

MTBF and Availability

DataDirect[™]
NETWORKS



- **MTBF is established by running a large drive sample**
 - 1000 drives running for 1000 hours without failure asserts a specification of 1×10^6 hours
- **Array MTBF=Drive MTBF/array size**
- **Availability=MTBF/(MTBF+MTTR)**
- **The goal in a large system is to reduce MTTR to a minimum to reduce data vulnerability**
- **Individual failure events must never affect data availability or performance**

MTTDL and Redundancy

DataDirect[™]
NETWORKS



- **MTTDL relates to the number of redundant elements in a group**
- **In a RAID 6 with dual parity**
 - A single drive failure results in a redundant system
 - A dual drive failure results in continued data availability
 - A triple drive failure results in data loss

What constitutes a failure?

- **Hard failures include**

- Head crashes
- Bearing wear
- Motor failure
- Electronic hardware failure (ASC\ASCQ 04)

- **Soft failures must include**

- Rereads
- Dynamic block reallocation
- Complete sector loss
- Data corruption
- Data recovery timeouts in excess of 20s

Agenda



- **Balanced I/O Systems**
- **MTBF and Availability**
- **Recovery Operations**
- **Parallel Storage Systems**
- **Conclusion**

Drive Error Recovery



- **Enterprise drives have less than 100 recovery steps**
- **SATA drives have over 200 recovery steps**
- **SATA recovery can range to 30 seconds**
 - Read and write of non-user data
 - Vary read amplifier characteristics
 - Re-read at +/- 6% of track width
 - Re-read at +/- 12% of track width
 - Adjustment of ECC parameters

Agenda



- **Balanced I/O Systems**
- **MTBF and Availability**
- **Recovery Operations**
- **Parallel Storage Systems**
- **Conclusion**

Parallel Storage Goals


DataDirect[™]
NETWORKS



Low Latency - High Performance, Silicon Based Storage Appliance

- Parallel Access For Hosts
- Parallel Access To A Large Number Of Disk Drives
- True Performance Aggregation
- Reliability From A Parallel Pool
- Quality Of Service
- Scalability
- Drive Error Recovery In Real Time
- True State Machine Control
 - 10 Virtex 4 FPGAs, 16 Intel embedded processors, 8 Data FPGAs

An Implementation of Parallelism w/ Double Parity RAID Protection



Single Enclosure Storage Tiering

DataDirect[™]
NETWORKS



Put data on the Most Effective Medium!




SATA
Cost Effective
Capacity



SAS
High Performance

S2A9900 Capacity

DataDirect[™]
NETWORKS



- Five 60-Slot Enclosures
- Two Connections per Enclosure:
300 Disks
- 300TB SATA using 1TB Drives
- 135TB SAS using 450GB Drives

- Ten 60-Slot Enclosures
- Two Connections per Enclosure:
600 Disks
- 600TB SATA using 1TB Drives
- 270TB SAS using 450GB Drives

- Twenty 60-Slot Enclosures
- Two Connections per Enclosure:
1200 Disks
- 1.2PB SATA using 1TB Drives
- 540TB SAS using 450GB Drives

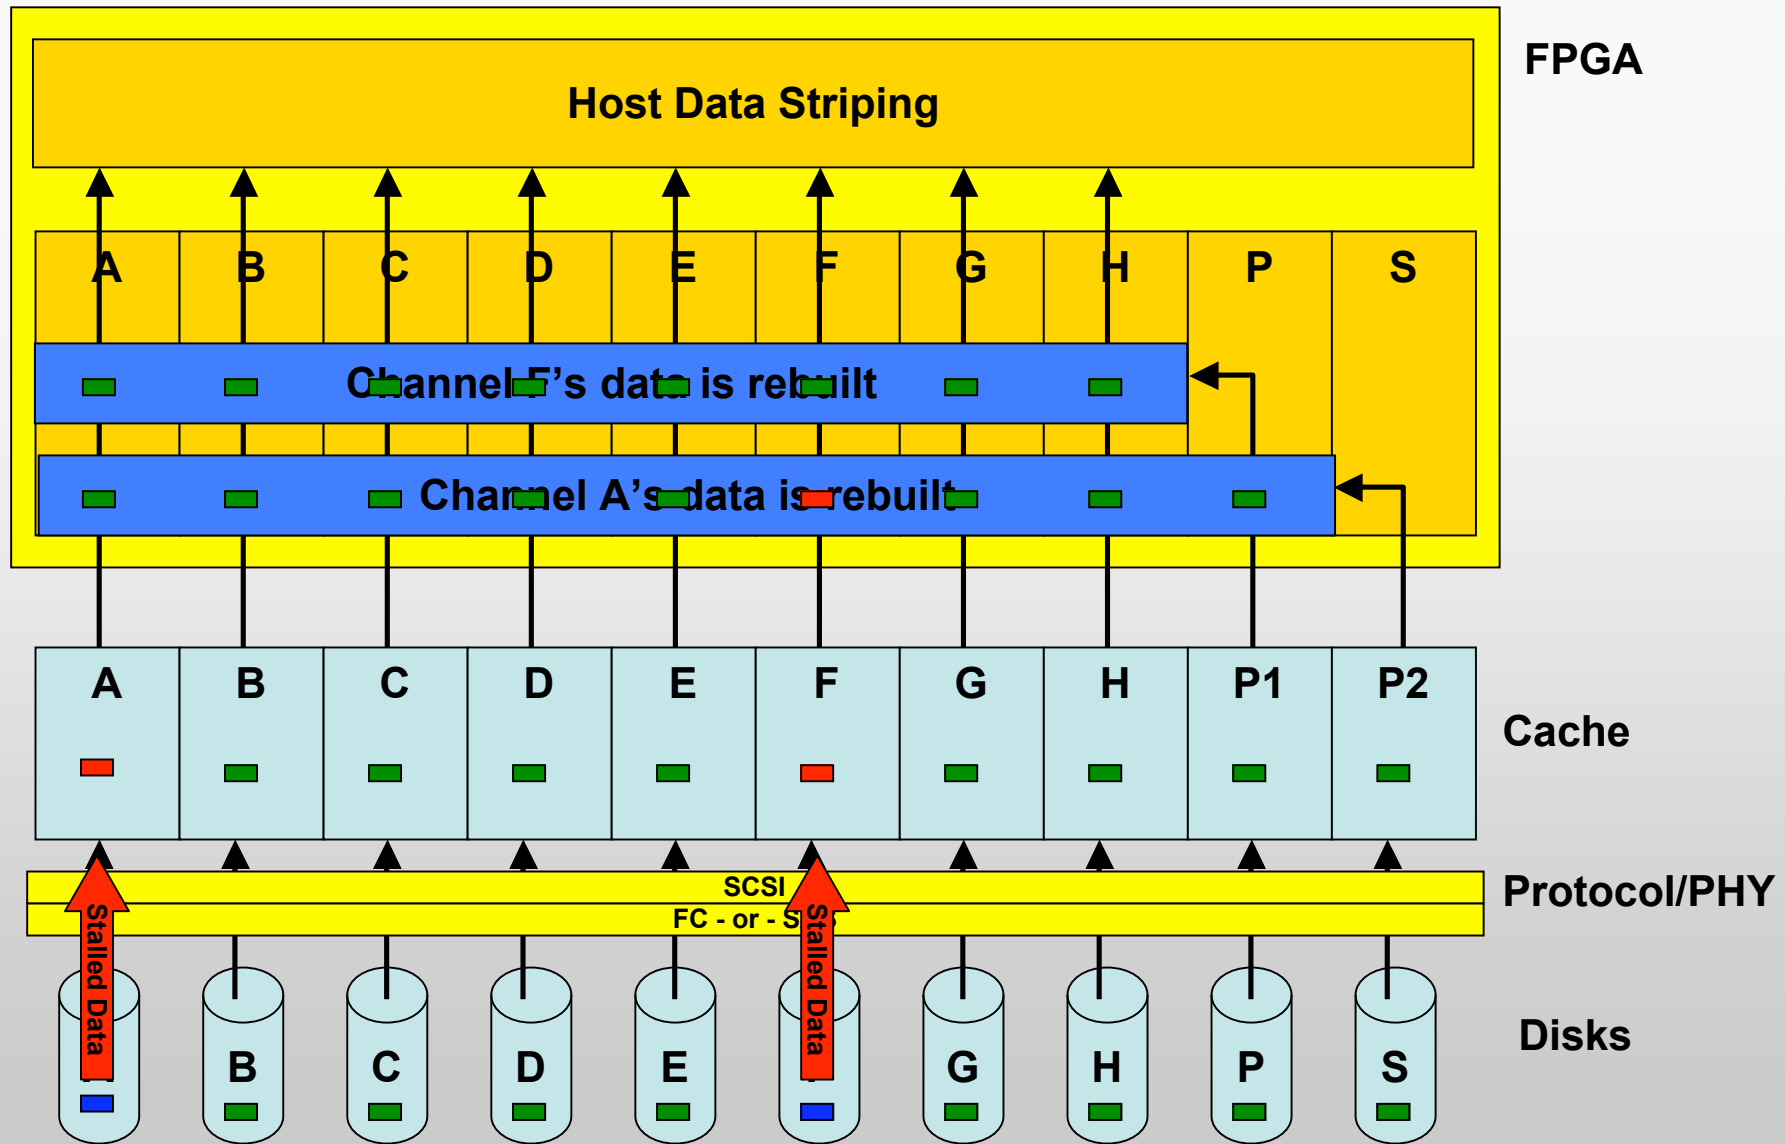
Quality of Service

DataDirect[™]
NETWORKS



- **S2A always reads (and writes) to all members of a RAID group**
- **S2A calculates parity on both Writes and Reads**
- **FPGA designed to generate host data with missing elements**
- **If a single member of RAID group is slowed by internal error recovery S2A can still provide host data at a high level of QOS**
- **All data passes through a Parallel Data Recovery Engine which recovers stalled or missing data**

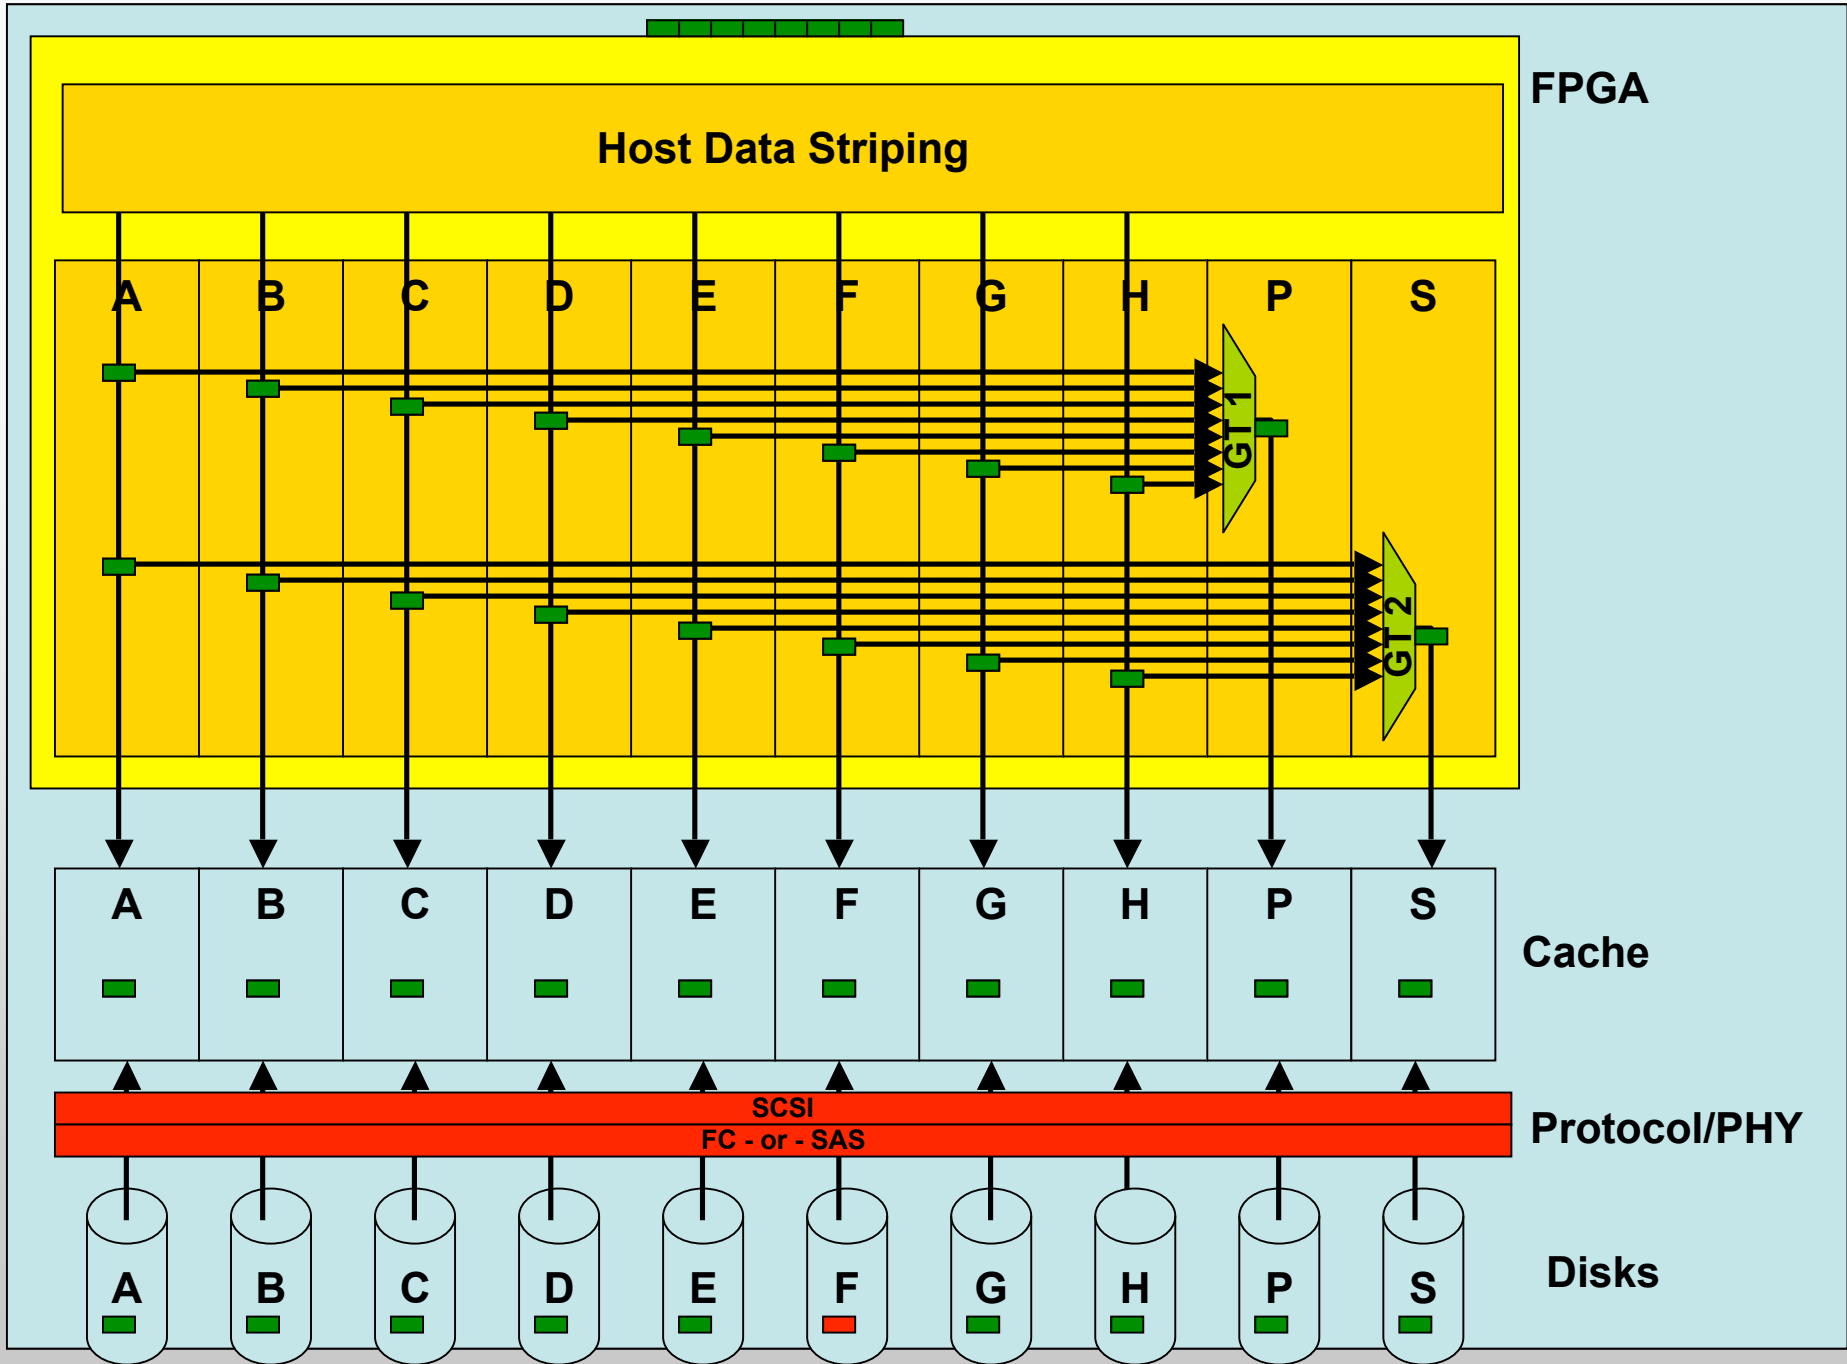
Quality of service



Data Corruption Error Handling

Note that the Cache and Disks have not corrected the data corruption.

We will need to rebuild the data into the cache and flush the data back to the disk in order to repair the problem fully.



FPGA

Host Data Striping

A B C D E F G H P S

GT1

GT2

A B C D E F G H P S

Cache

SCSI
FC - or - SAS

Protocol/PHY

A B C D E F G H P S

Disks

SATAssure Data Integrity

DataDirect[™]
NETWORKS



- **The Parallel Data Recovery Engine allows data reconstruction and integrity checking**
 - S2A hardware enables SATAssure software to verify all data read from the disks
 - S2A hardware allows SATAssure to send hosts “fixed” data (**data integrity is assured**)
 - S2A hardware enables SATAssure to correct data on the disk for future accesses (**self-healing array**)
 - Multiple levels of disk recovery attempted before failing drives (**replace fewer drives**)
 - S2A controller journaling allows partial rebuilds (**less time in degraded mode**)

Worst Case Recovery

DataDirect[™]
NETWORKS



- **Disks can become completely unresponsive to all commands**
 - The internal OS can enter a loop that does not enable external commands
 - A power cycle always recovers the drive
 - S2A 9900 automatically power cycles a drive in place
- **SAS drives can be issued LLF in place**
 - Platters are rechecked for integrity and the sectors are rewritten
- **Drives that issue SMART warnings or grow defects at an increasing rate are copied to spare drives**

- **Does data checking and recovery have an impact on performance?**

On S2A: Not

Agenda



- **Balanced I/O Systems**
- **MTBF and Availability**
- **Recovery Operations**
- **Parallel Storage Systems**
- **Conclusion**

Conclusion

- **Bit error rates and drive error recovery mechanisms are a statistical reality**
- **Large clusters must maintain a very high data rate to minimize the I/O cycle and have predictable I/O times**
- **Storage systems must execute self test and repair to minimize human intervention in the machine room**
- **If human intervention is required every possible automation assist must be employed**

==> NTF drives reduced from >65% to <20% on product that can support RAID 6 error recovery

Future Technology

DataDirect[™]
NETWORKS



- **Systems must be kept as small and power efficient as possible**
- **SSD technology must be utilized in conjunction with rotating media**
- **AI must be used to simplify management**
- **File system service must be a part of the storage system**
- **The storage system must be capable of data analysis**

DataDirect[™]
N E T W O R K S



Thank You

Toine Beckers
tbeckers@datadirectnet.com

Number One Oil Producer on Extreme Storage

DataDirect[™]
NETWORKS



When you're the World's #1 Oil Producer,
You Need Extreme Storage. S2A Delivers.

أرامكو السعودية
Saudi Aramco



Saudi Aramco supplies **10% of all the oil consumed in the world**. They needed highly advanced storage technology to meet their **Extreme Seismic Exploration** needs.

With over 50 S2A systems from IBM, Saudi Aramco is equipped to accelerate time to oil.