

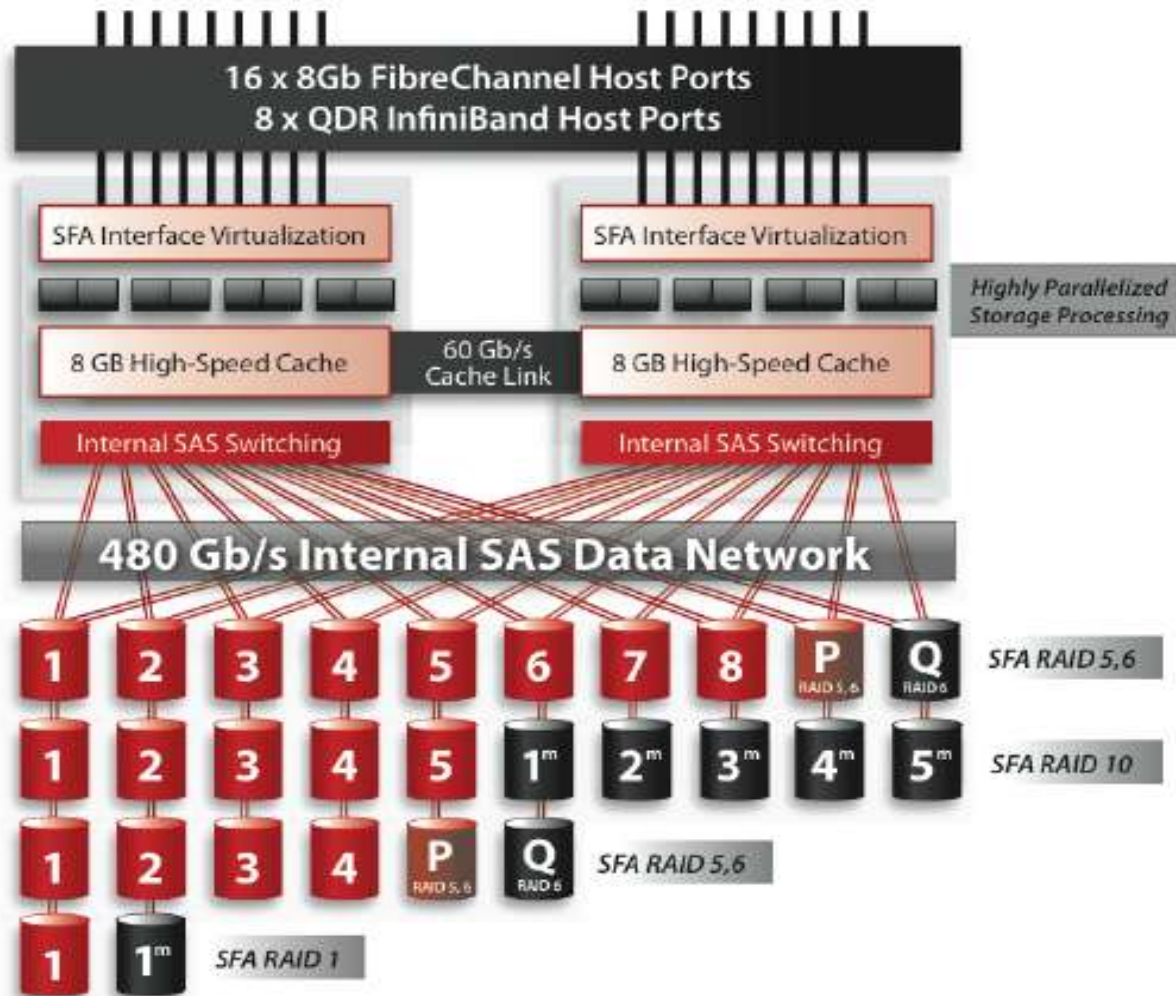
Storage IO Challenges

Ghislain de Jacquelot

Technical Director HPC EMEA

DataDirect[™]
NETWORKS
INFORMATION IN MOTION[™]

Overview of the SFA10000



Optimized Drive Support



Leading Capacity and Cost-Optimized Bandwidth



Balanced Mix of IOPS, Capacity & Bandwidth



Delivering Unrivaled IOPS for Transactional Applications

Hardware IO challenges

Provide simultaneously

- High-Availability
- Reliability of data
- High Bandwidth
- High IO rates
- Undegraded performance
- Mixing cheap components and fast devices

- **Read error rates**
- **Rebuild performance**
- **Mix of small & large IO**
- **In-Store compute**




Dealing with read error rate

Disks: binary devices ?

- **Either new or full**
- **Working or dead ?**
- **Stores 0s and 1s ?**

Reliability & Performance Challenges

Error Rate Constant at 1 in 10^{15} bits read



- Disks are >100 times denser than a decade ago, but media error rate is constant
- ***More Media + More Media Errors per Drive***


Defense-In-Depth Solutions

- **DDN's SFA Protects, Detects, Corrects & Enables Performance in Multiple Ways**
 - Bit Error Rates Get Eliminated Through Parity
 - Checksums & Parity Calculation on Read Options
 - Foundation for T10-Dif (like) Protection Without Perf Loss
 - Background Parity Scrubbing (cycle data at rest)
 - Robust cache protection eliminates corruption on write

Much of this technology also serves as basis for future real-time capabilities.

Defense In Depth: SATAssure

300 x 1TB SATA Drives, 30 x 8+2, Write Throughput,
Varying Block Size, 128K Chunk Size, 16 x FC-8



6Gb SS7000 Enclosure Increases Performance 10%

Dealing with rebuilds

DataDirectTM
NETWORKS
INFORMATION IN MOTIONTM

10PFlop Deployment Scenario

- **50 PB Useable Storage System**
 - 15,000 4TB Drives = 1,500 LUNs
 - 1TB/s Storage Cluster
 - Drive AFR = 3%
 - AFR is observed data from all DDN installation base
 - **1.2 drive failures per day**
 - Without fast rebuilds – file system is always rebuilding
- **Without Performance Protection...**
 - A striped I/O will only run as fast as the slowest LUN which will be the LUN being rebuilt

Reconstruction Comparison


	SFA Software	Software RAID
Capacity Utilized	60%	60%
Rebuild Speed	40MB/s	25MB/s
MTTR	28 hours (all blocks)	26 hours (data only)
Performance Impact	10%	50% When Optimized (Currently 63%)
Striped IO Perf In Reconstruction	820GB/s	455GB/s (best case) 233GB/s (unoptimized)

1.2 drive failures per day with rebuild taking a full day:

The file system is always slow


SFA10000

- **Maximum Data Availability & Performance**
 - Direct cabling avoids daisy chaining & latency
 - Data is striped across channels/enclosures for maximum bandwidth to/from HDDs
- **Designed To Withstand 20% Loss without Service Impact**
 - Drive Channels are RAIDed 8+2
 - Drive Enclosures are RAIDed 8+2
- **Oversized 60GB/s Backend**
 - Preserves performance even during rebuilds
 - Supports data protection and management IO



Ahmdal's Law & Striped File I/O

No More Than A 10% Performance Drop For Rebuilds, Compared To As Much As 66% With Other Systems



Mixing large and small IO

DataDirect[™]
NETWORKS
INFORMATION IN MOTION[™]

ReACT technology

Bringing out the performance of SFA10K


Write-Back Cache $\hat{=}$ Bottleneck in Mirror & Interconnect

Aligned IO


- Written through cache and striped directly to disk
- Relieving the ICL and avoiding parity overhead

Non-Aligned IO

- Mirrored Write-Back Cache
- Limited to 6GB/sec (QDR ICL x2)



Performance Efficiency




With ReACT, partially striped I/O is given more cache headroom to achieve full stripe & optimize applications. SFA Algorithms ensure max. performance per array, per cache GB.

In-Store compute

DataDirect[™]
NETWORKS
INFORMATION IN MOTION[™]

In-Store Compute



- Announced In 2009
- Optimized For Parallel File Systems
- Supports File I/O, Pre and Post Processing

IO Path Acceleration

IO Paths

Latency

Traditional



SFA 10000E



Applications Efficiently Place Data Directly Into SFA Memory



Elimination of Protocol Conversion Reduces Latency, Improves IOPS Performance.

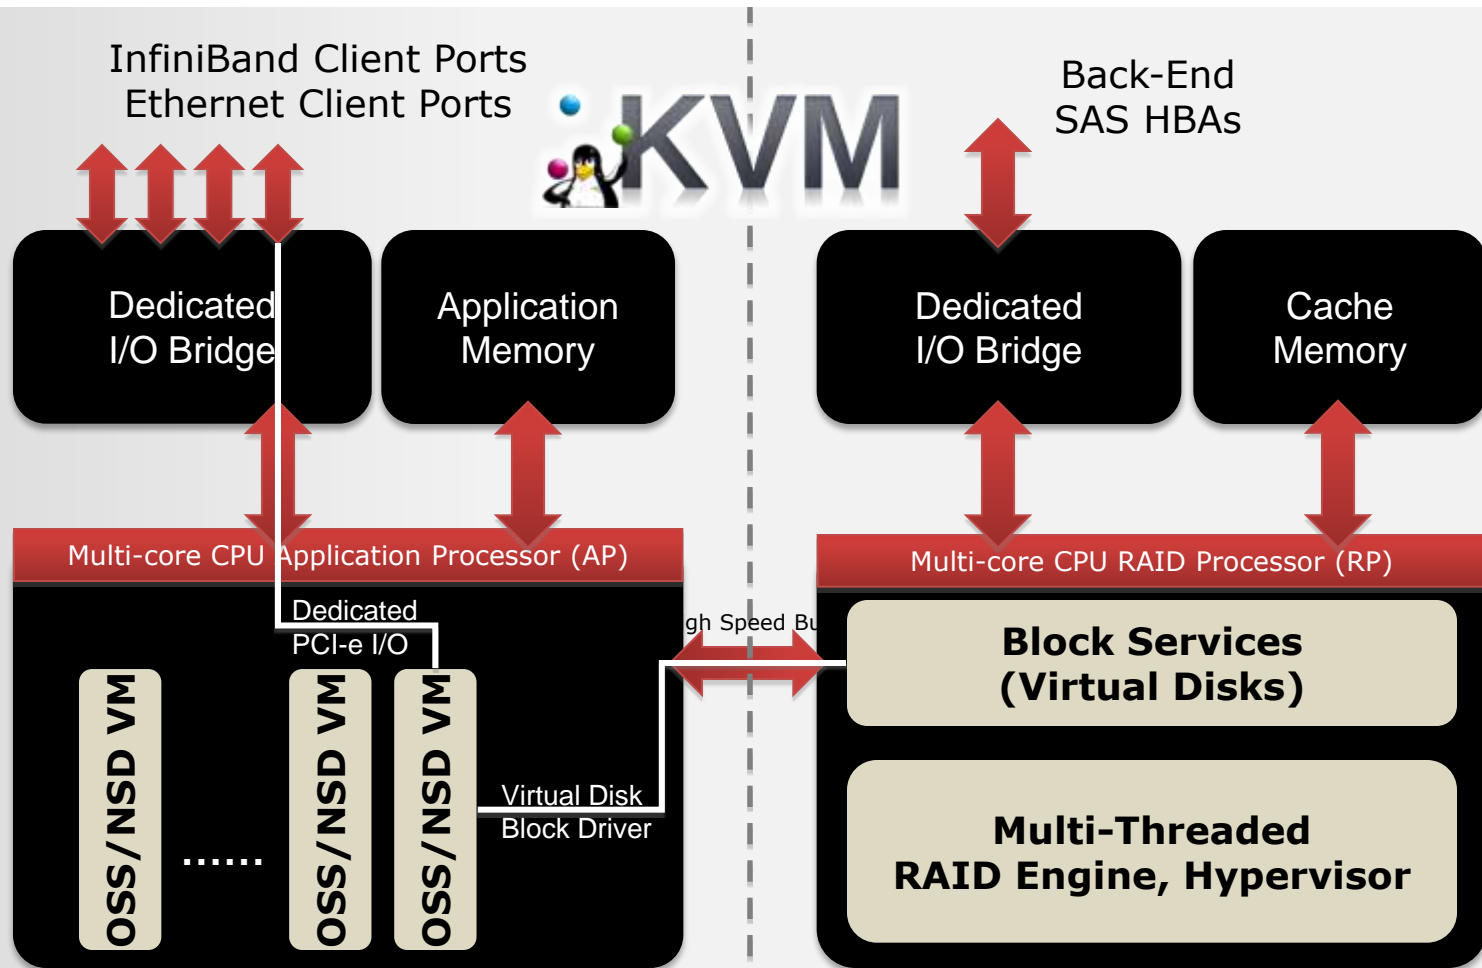
Storage Fusion Architecture shortens the IO path from the application to storage, reducing latency and increasing IOPS performance.

Appliances, We've Got 'Em

- **Embed Apps Within the Storage Array**
 - Reduces complexity, infrastructure and administration
 - Reduces cost as well as lower operational cost
 - Increases performance – access to disk is direct
- **Appliance-ize File Systems**
 - Eliminate external networking and filer heads
 - Reduce file system licenses
 - Virtualization Enables Multi-Port Storage Servers
 - Increases performance, particularly for small I/O



In-Store Compute Virtualization



GPFS Embedded on SFA


FAST

Traditional block storage, direct attached to Parallel FS Servers

Parallel File System
Clients




IB or 10 Gig-E



File System
• Networking
• Server Infrastructure
• NSD Licenses

IB



FASTER, SIMPLER

Servers embedded in 10K, improving performance and reducing network and servers

Parallel File System
Clients



Thank You

Questions?

DataDirect[™]
NETWORKS
INFORMATION IN MOTION[™]