Accelerating NoSQL on ioMemory

November 2013

Fusion-io at a Glance

- **Founded:** December 2005 to solve the **Data Supply Problem**
- **Operations:** Salt Lake City (HQ), San Jose and Denver
- **Employees:** ~750
- **Results To Date:**
  - More than 3,500 customers across multiple verticals
  - Significant deployments at key accounts
  - Customers achieving > 10x increase in application performance
  - OEM relationships with Cisco, Dell, HP, IBM and more
- **IPO:** June 9, 2011, Symbol: FIO

Fusion-io First Mover Milestones

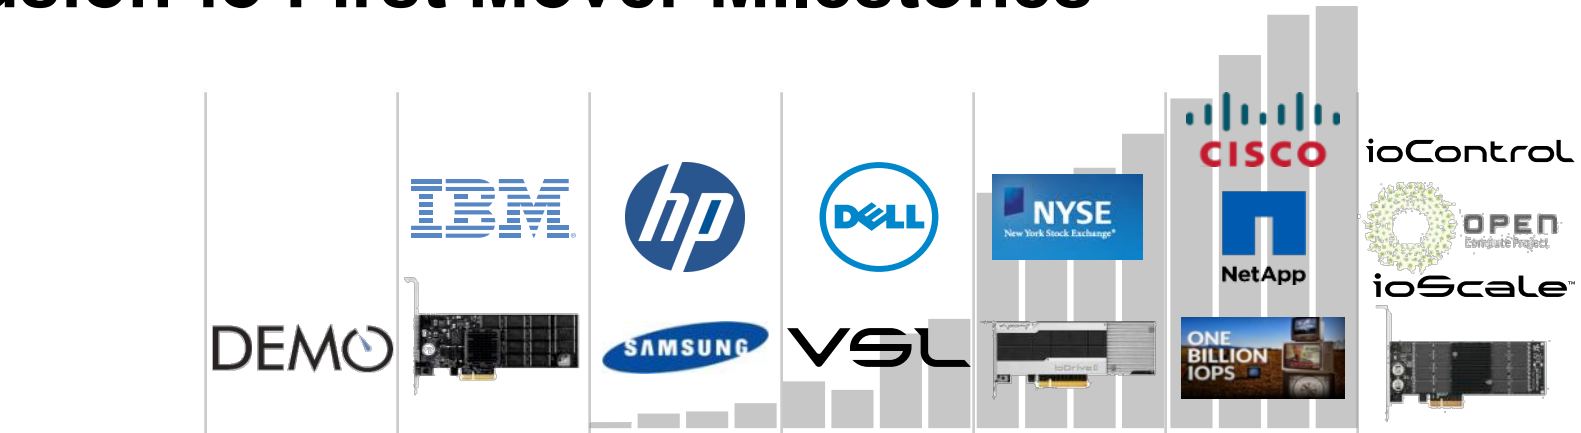

2006

Mission to consolidate memory and storage

2007

ioMemory technology unveiled

2008

First products launched

1 million IOPS IBM Quicksilver

Dell strategic investment

2009

HP OEMs products

IBM OEMs products

Samsung strategic investment

2010

Dell OEMs products

VSL introduced

2011

IPO on NYSE

ioTurbine acquired

ioDrive2 announced

Supermicro OEMs products

2012

1 Billion IOPS

2,500 customers

>120 channel and alliance partners

ioFX announced

ioMemory SDK introduced

Cisco OEMs products

2013

ioScale announced at Open Compute Summit

NexGen acquired

Fusion-io Accelerates

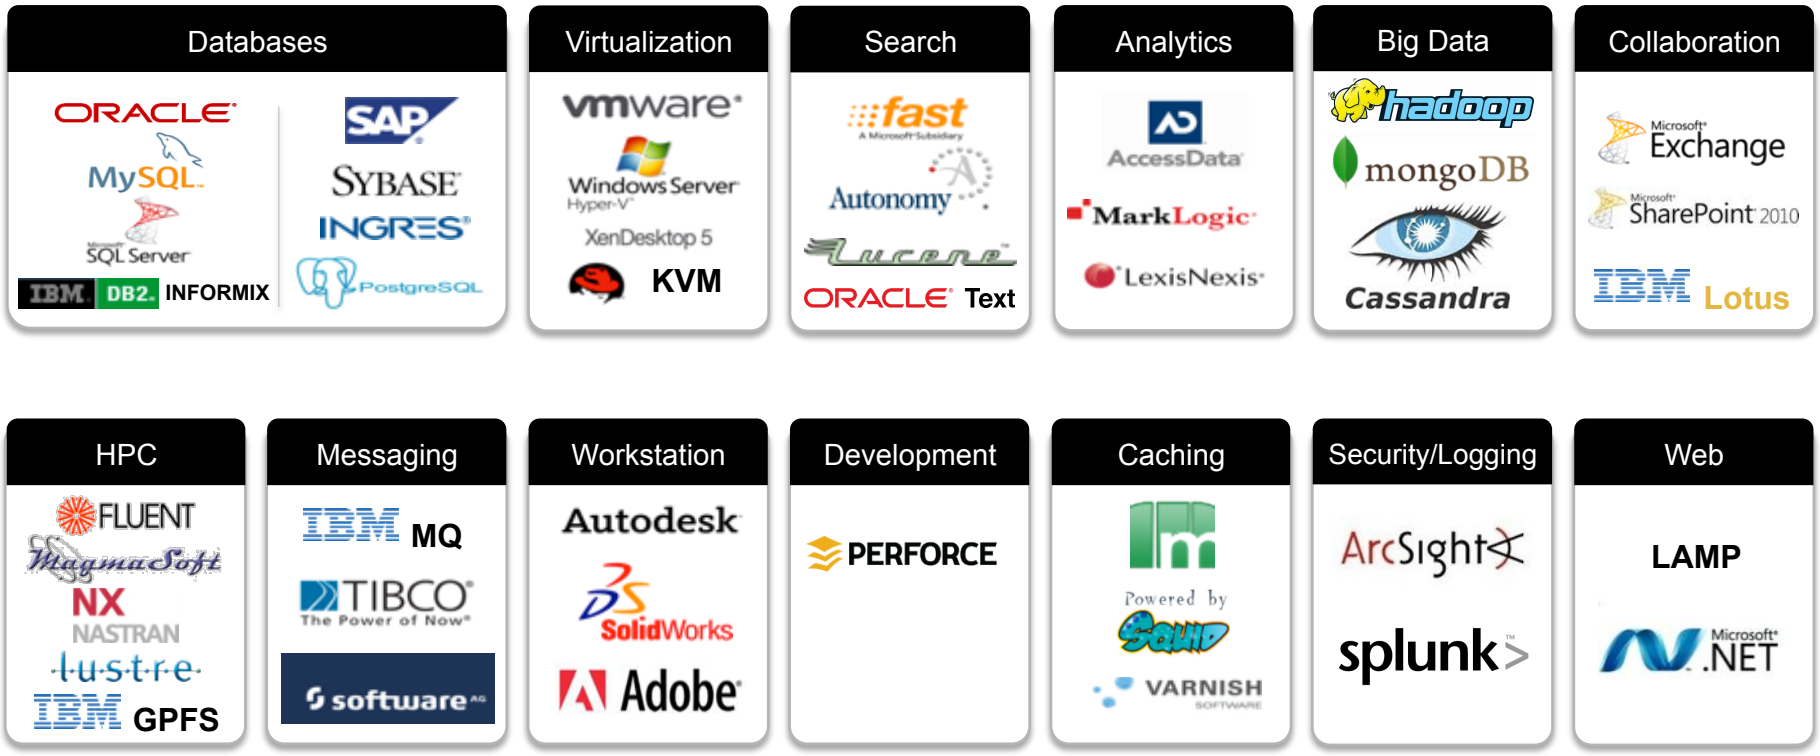

Data Supply Problem

- Processing performance doubles every 18 months
- But storage performance has not kept up

37% OF SERVERS ARE UNDERUTILIZED[1]

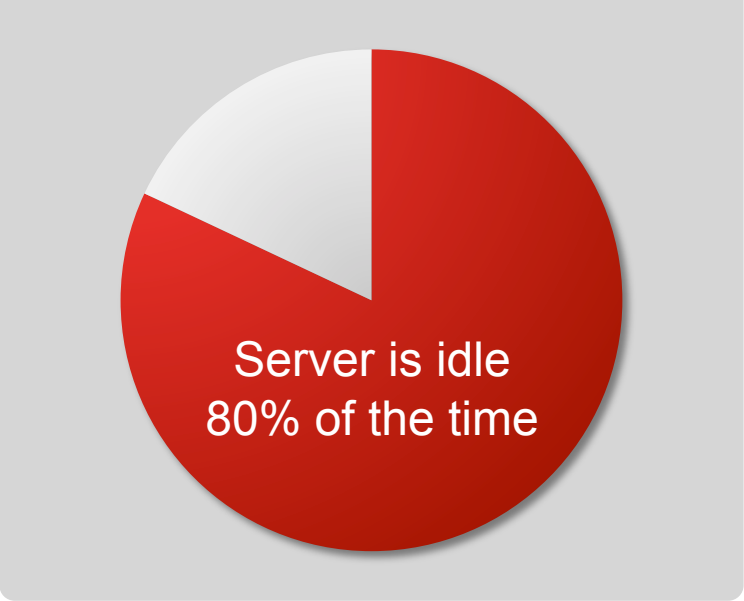

PERFORMANCE GAP CONTINUES TO GROW

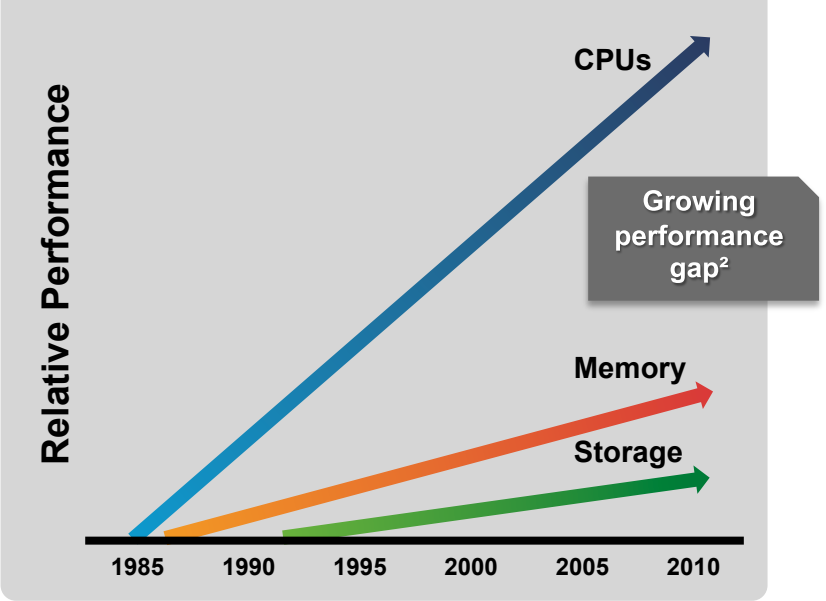

[1] *Source: IDC's Server Workloads 2010, July 2010*

[2] *Source: Taming the Power Hungry Data Center, Fusion-io White Paper*

Traditional Centralized Architecture

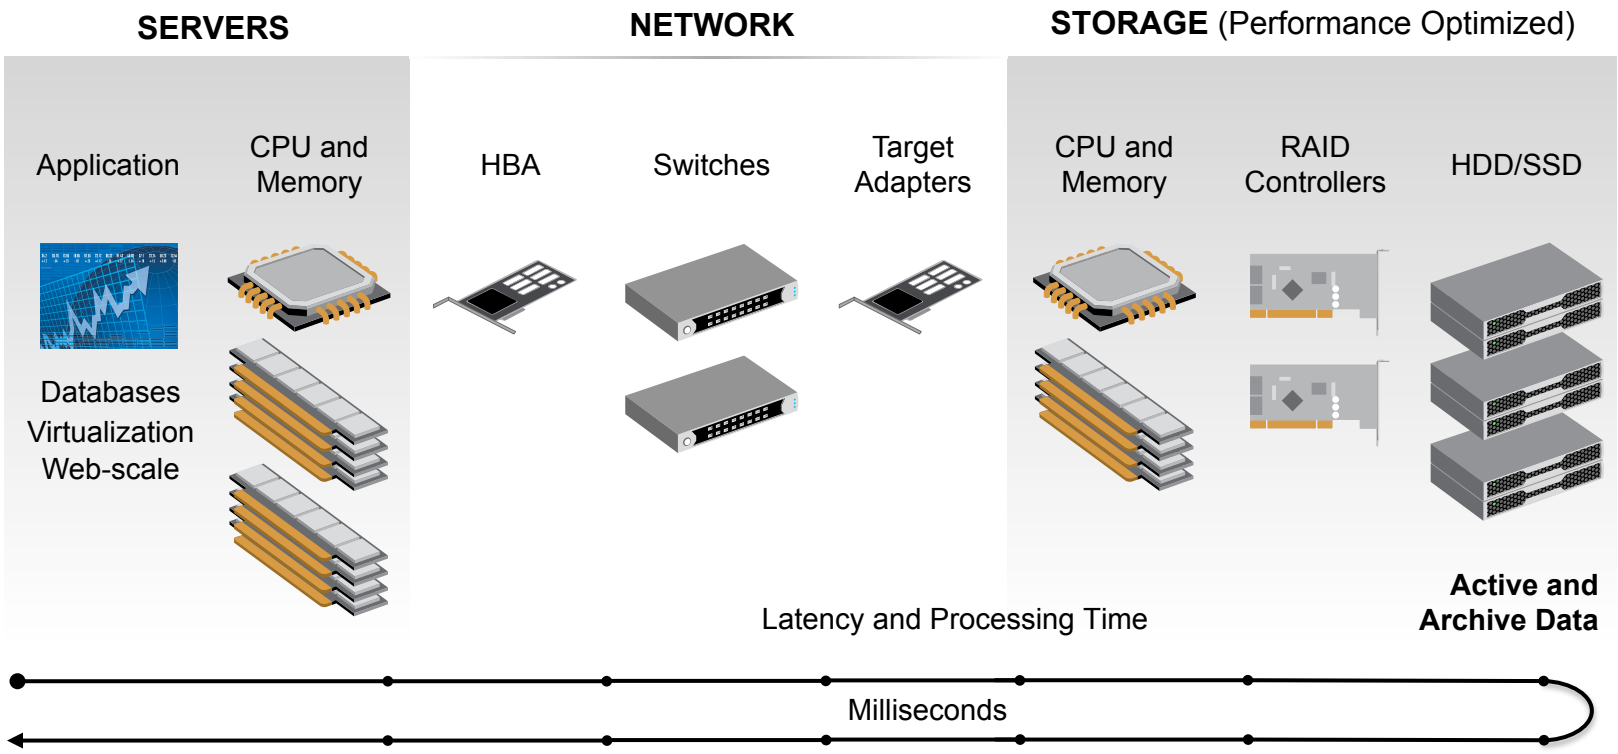

Shared Data Decentralization

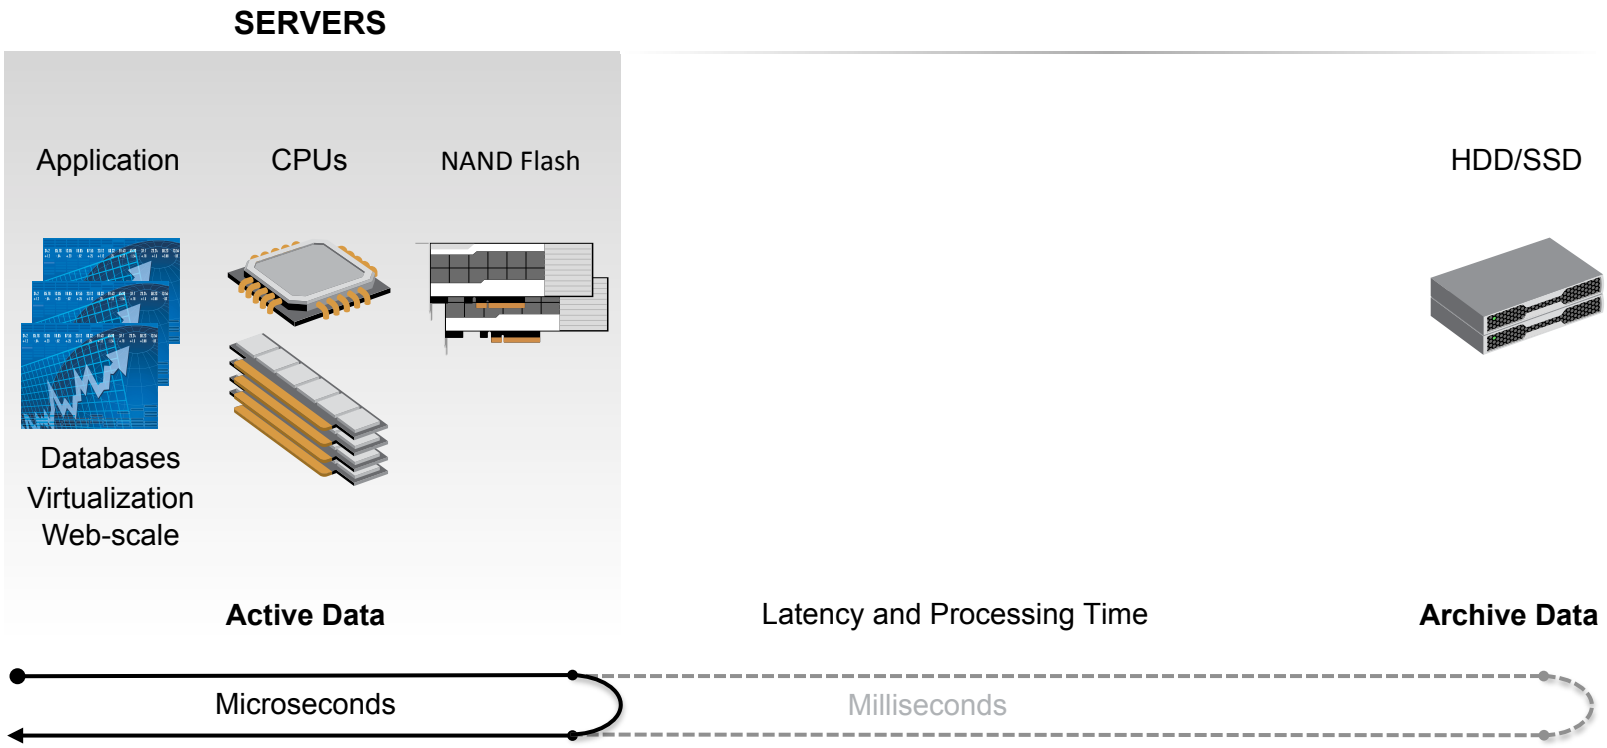

Cut-through Architecture and VSL

- Sophisticated architecture
  - maximum performance
- Intelligent software
  - advanced features

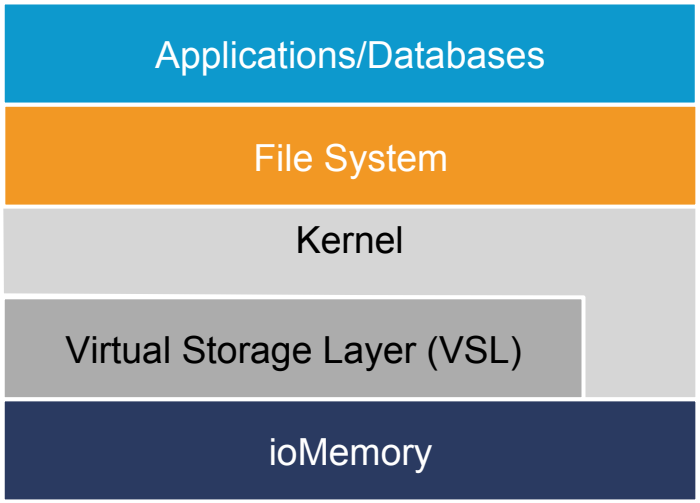

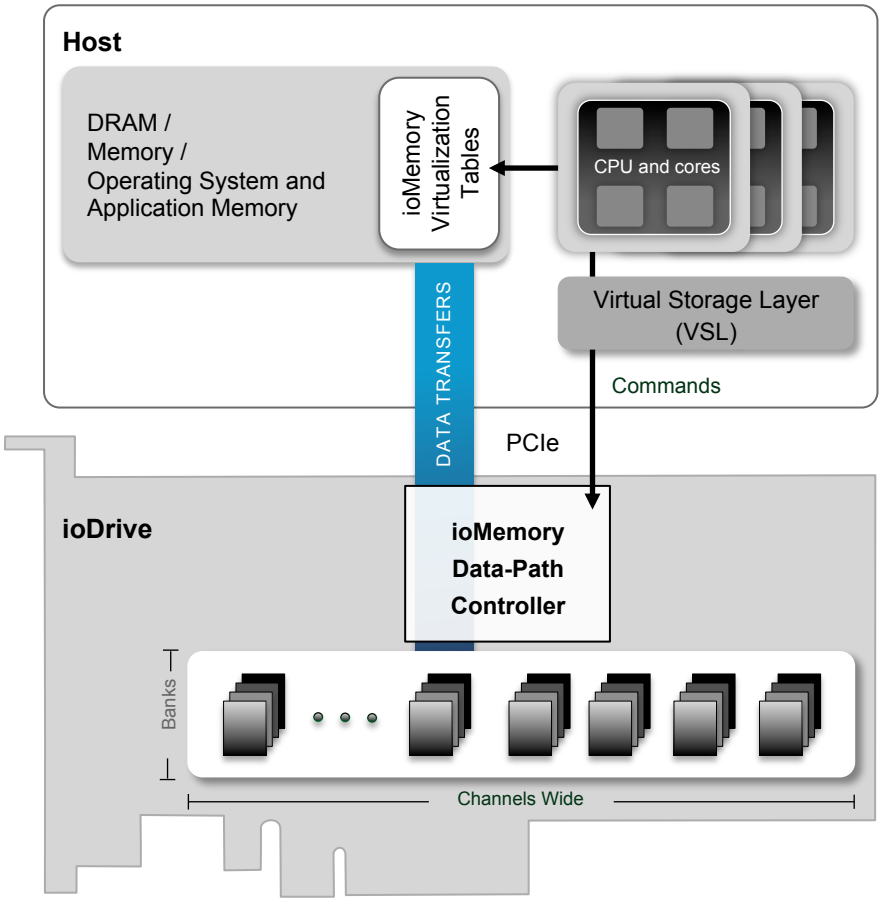

Flash Architectures

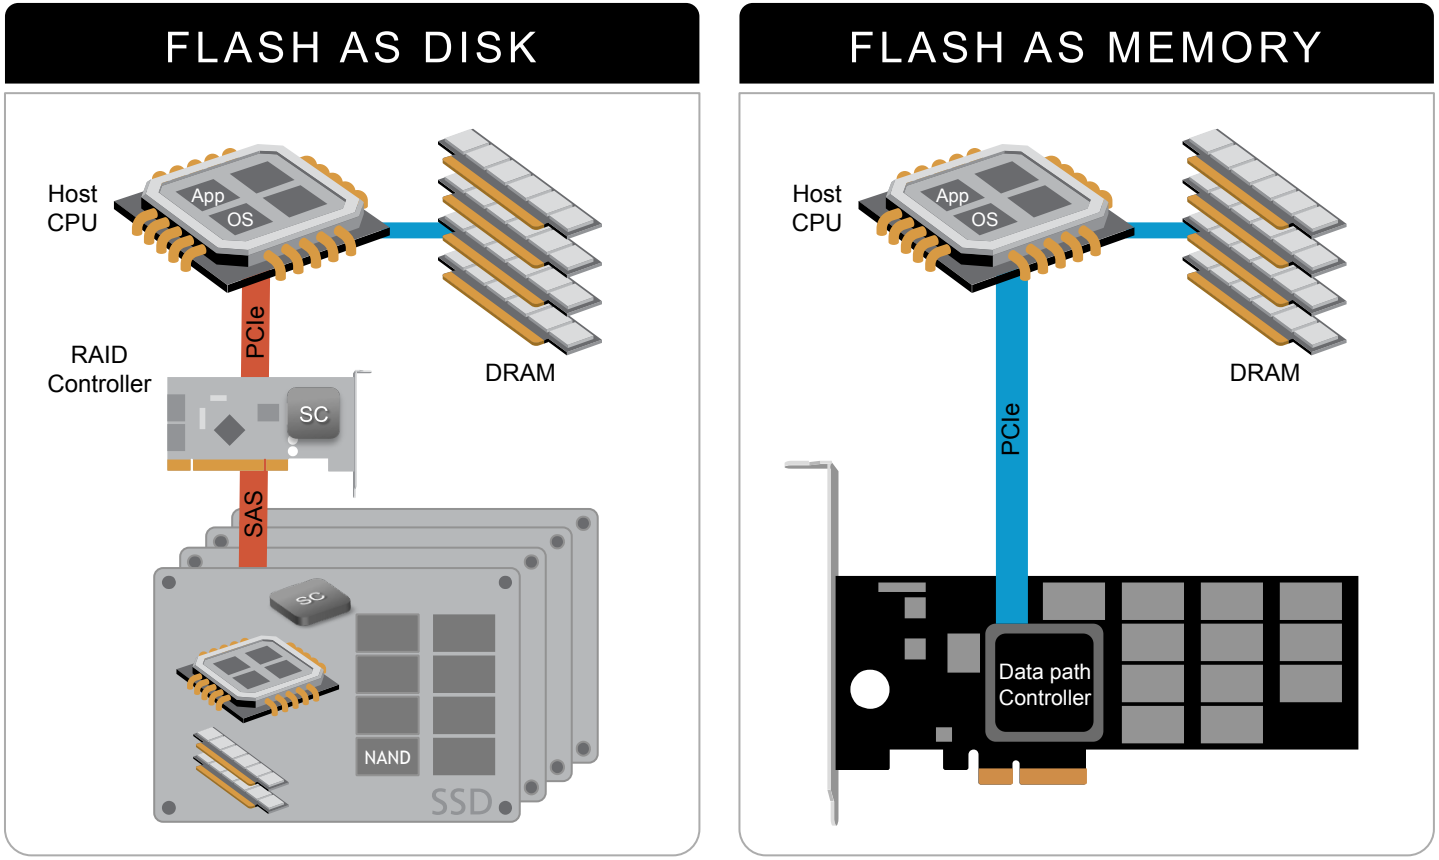


FUSION-IO®
Product Introduction
November 2013

Fusion-io Product Portfolio

DIRECT	CACHING	SHARED
Max Acceleration • Closest to CPU • Fastest deployment • Scale with servers	**Max Interoperability** • Virtualize more • Consolidate server • Offload SAN	**Max Control** • Allocate across servers • Scale independently • Hybrid and QoS options
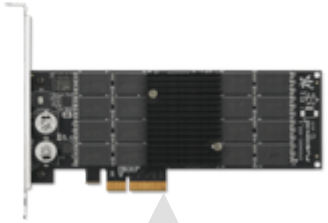		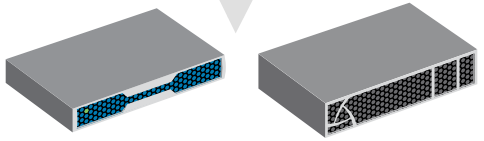
ioMemory®	ioTurbine™	ioN | DATA ACCELERATOR ioControl | HYBRID STORAGE

Direct Acceleration

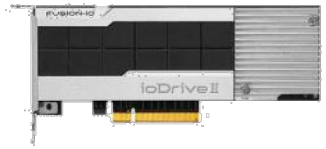

ioDrive® II

Up to 3.0TB of capacity

ioDrive® II Duo

Up to 2.4TB of capacity
per x8 PCI Express slot

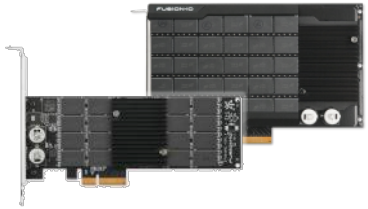

ioScale™

Up to 3.2TB of low-latency,
high-performance flash per PCI
Express slot

ioDrive® II
MEZZANINE

Up to 1.2TB for maximum
performance density

Caching Acceleration

ioTurbine™

- Purpose-built performance for virtualized applications and databases
- Unparalleled low latency performance
- Increase VM density and consolidate servers

ioCache®

- Turbo Boost virtualization
- Transform ioMemory into a powerful, easy-to-manage, intelligent cache
- Unleash the potential of virtualized systems

Shared Acceleration

ioN | DATA ACCELERATOR™

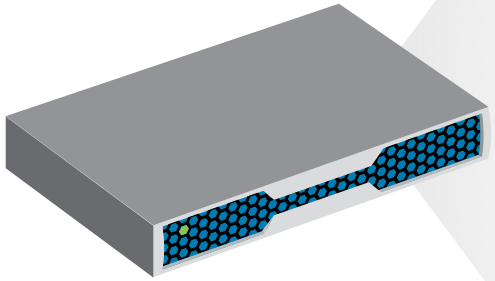

Software

- 25-40x performance
- 100µs response time for cluster applications

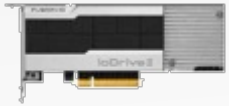

ioMemory

- PCIe performance
- Adaptive Flashback reliability

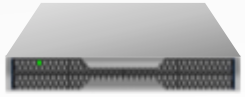

Platforms

- SAN Connected (FC, iSCSI, Infiniband)
- Qualified servers

25x DATABASE PERFORMANCE

40x INDEX REBUILD

35x XPLORE QUERY TIMES

Hybrid Acceleration

ioControl™ | HYBRID STORAGE

iSCSI Hybrid SAN

- Integrated appliance using ioMemory and disk
- Shared ioMemory performance for $1-$5 per GB

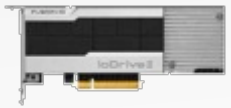

ioMemory

- Proven, tested, reliable and fast
- Thousands of customers

Mission Critical
Business Critical
Non Critical

Performance Control

- Provision performance just like capacity
- Prioritize workloads

10x PERFORMANCE

2x CAPACITY

1/3 DATA CENTER FOOTPRINT

FUSION-IO

Flash Optimization

ioSphere®

- Data Center ioMemory Management from a single interface
- Real-time monitoring and management

VSL™

- Virtualizes flash memory
- Direct ioMemory access
- Remove bottlenecks of disk-era RAID controllers and storage protocols


FUSION-IO
Big Data - NoSQL
November 2013

FLASH makes Big Data more Efficient

In Big Data, implementing flash is not just about raw performance – its also about architectural efficiency.

Architectures built solely on DRAM for performance and spinning disk for capacity will suffer from the inefficiencies of both media.

Fusion-io can entirely replace spinning disk storage and heavily reduce the DRAM footprint.

Fusion-io can also complement disk storage as a cache layer

Either strategy can **improve efficiency**, **increase density** and **reduce operational costs.**

Big data “building blocks”

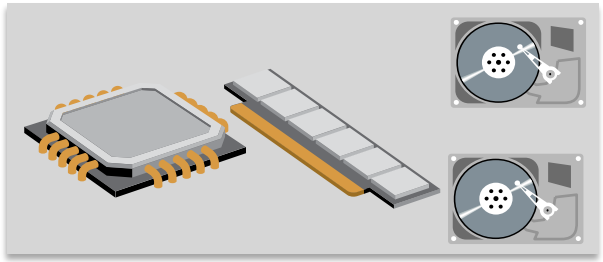

Tend to make heavy use of DRAM to accelerate reads.
Tend to serialize write I/O
Need “working set” < “Total DRAM” to maintain performance

DRAM dictates NoSQL Scaling

- As the working set increases:
  - DRAM pricing and capacity quickly become an obstacle to efficient scaling.
    - Low density limits DRAM scaling to a few hundred GB per server
    - Pricing increases substantially at higher DRAM capacity points

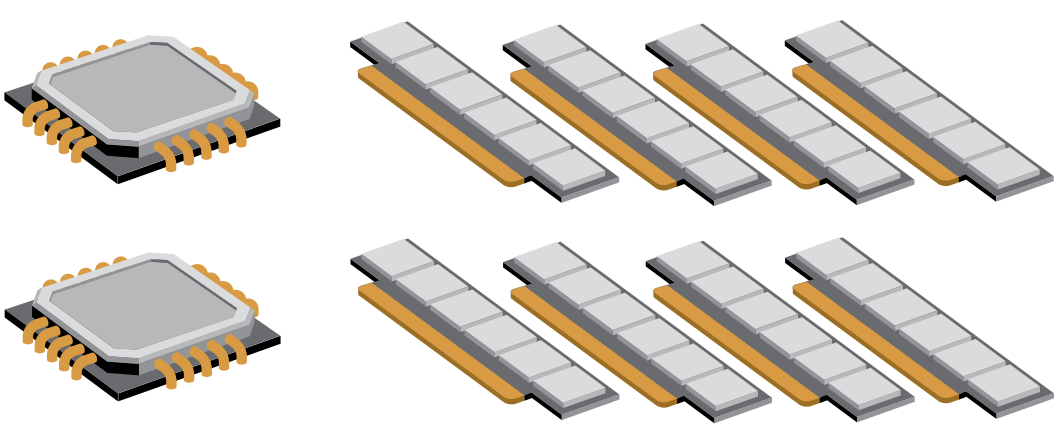

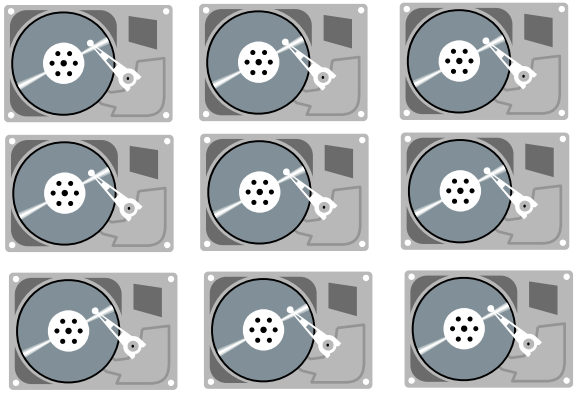

Consider..

- Traditional HDD and SSD systems meet performance through scale out.
  - The DRAM price curve..
    - Quickly becomes cost effective to buy another server and use cheaper, low density RAM..
  - Therefore adding more servers, cores, spindles.
  - This reduces the efficiency of the datacenter and increases solution costs.
  - Power, cooling and rack space at a premium.

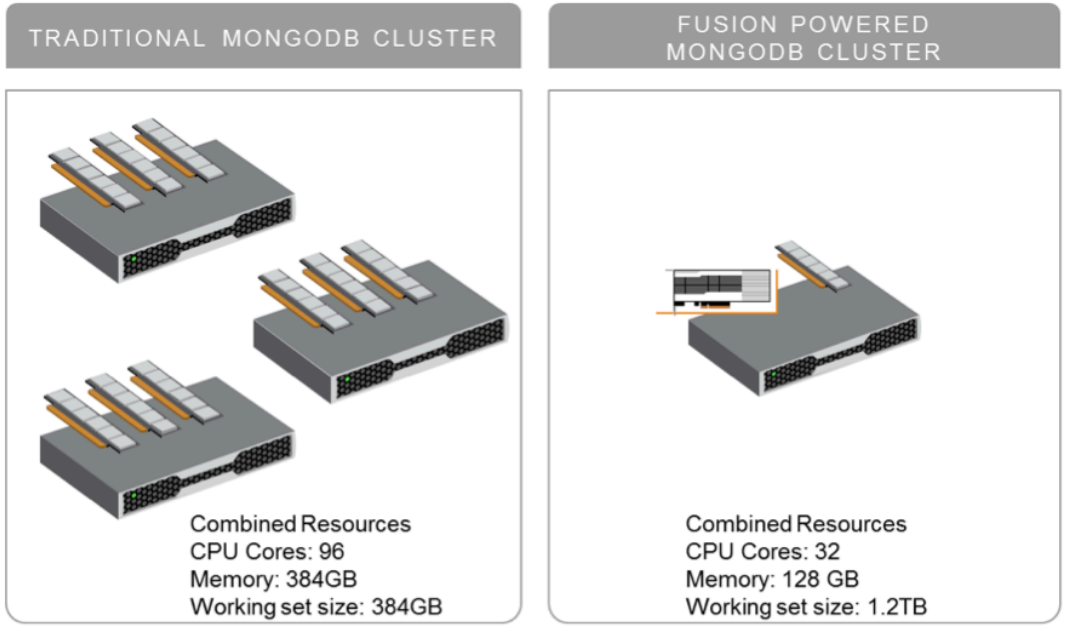

- Above diagram shows 3 Servers for 384GB working set compared to 1 server with a 1.2TB working set.
- Consolidation saved 1,100 Watts in power and a further 1,100 Watts in cooling
- 66% reduction in rack space (6u to 2u).

All-Flash Storage

- Fusion ioMemory as primary storage.
  - Readily available at 10x the capacity of DRAM *per PCI slot*
  - At approx. 1/10th cost of DRAM
  - Requests served directly from persistent flash memory
    - Read response times across the entire database now similar to DRAM “cache hits”

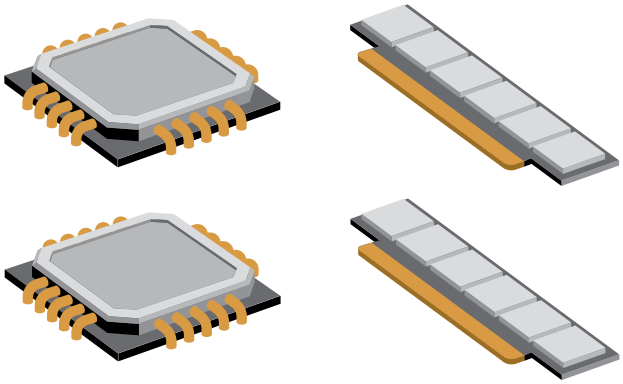

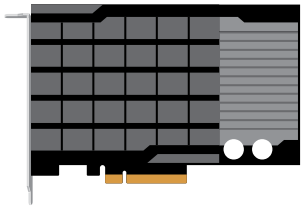

MongoDB response times

UP TO 40X IMPROVEMENT IN MONGODB RESPONSE TIMES

We ran tests comparing a big data workload (specifically Yahoo! Cloud Serving Benchmark (YCSB)) on ioMemory versus 10 x 7,200 RPM HDD hard disks in a RAID 0 under the following workloads:

- Workload A: 50/50% read/write mix
- Workload B: 95/5% read/write mix
- Workload F: Updates
- Workload F: 50/50% read/read+modify+write mix

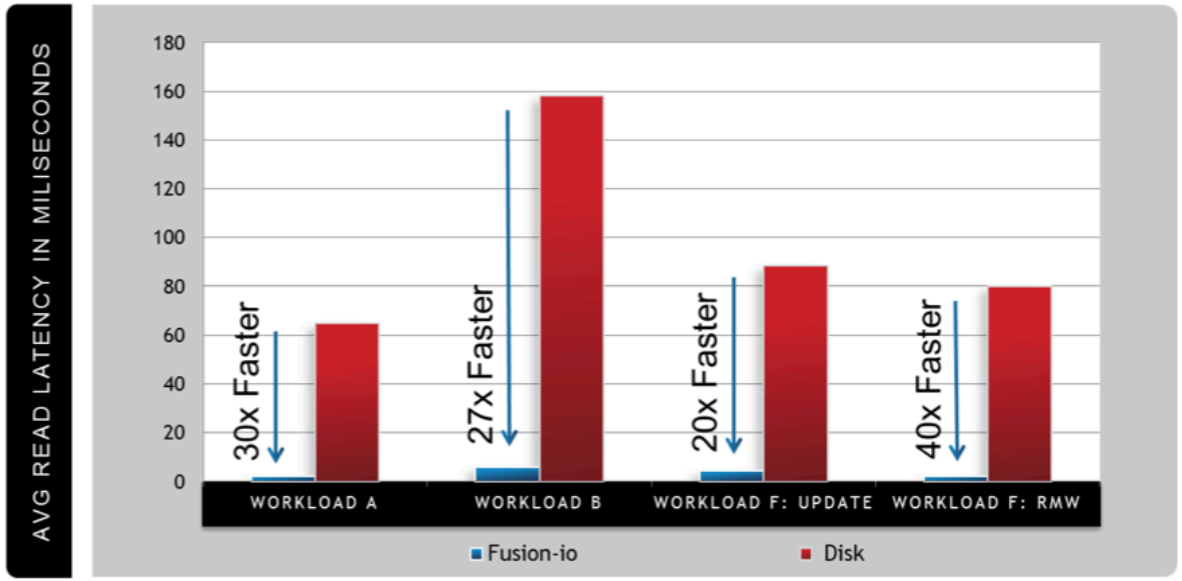

Conclusion

- 11-18x write performance improvement
- 20-40x read performance improvement
- Read latencies massively reduced
- 2.2KW saving in power and cooling
- 66% reduction in rack space
- No need for sharding of the databases

Spotify®

- Over 24 million active users
- Over 20 million songs available globally
- Over 6 million paying subscribers
- Over 1 billion playlists created
- Over $500 million paid to rights-holders
- Over 850 employees
- Over 250 developers

- Available in: 28 countries - USA, UK, Australia, New Zealand, Germany, Sweden, Finland, Norway, Denmark, France, Spain, Austria, Belgium, Switzerland, The Netherlands, Ireland, Luxembourg, Italy, Poland, Portugal, Mexico, Singapore, Hong Kong, Malaysia, Lithuania, Latvia, Estonia and Iceland.

- Over 24 clusters and quickly growing.
- Containing over 300 nodes
- Distributed over 4 data centers around the world
- Our main solution for scalable storage

Why Flash?

- “It changes everything, is a step change going from spinning disks to flash”.
- “Cassandra is page cache bound - flash moves scaling from memory to flash”.
- “Allows us to both consolidate and scale our clusters at the same time”.
- “Developers can focus on delivering products instead of optimizing for I/O”.

Why Fusion-io?

- "Why attach flash to a legacy platform"?
- "It turns out that it's easier to get installed".
- "Performance".

Early Results:

- 3-4x consolidation factor.
- 3-6x reduction in latency.
- Forcing SStables to memory not needed anymore.
- ROI so far is 2.2x
- Consolidation limited by Cassandra 1.1

"Spotify users expect fast results across all of their devices. Fusion ioMemory gives us the speed and scalability we need to grow our footprint worldwide with new services and scale our user base by the millions,"
Patrik Torstensson, Architect at Spotify

Summary

Scalable architectures need balanced components:
CPU, memory & low latency, persistent storage.

Fusion-io enables predictable, consistent high performance.

Increase workload density.

Reduce complexity.


Thank You
fusionio.com | SAME PLANET. DIFFERENT WORLD.