



Fusion ioDrive Sparks a Storage Revolution

Analyst: Michael Fisch

Management Summary

A hardware technology exists that will seriously challenge hard disk drives at the high end of enterprise storage. The technology is NAND flash, and the company pioneering its use as primary enterprise storage is Fusion-io. In the first quarter of 2008, it will release a product called *ioDrive* that delivers the IOPS performance of a large SAN array in the form factor of a PCI-e card. Capacities range from 80 to 640 GB. **Read on for details about how and why this product by Fusion-io will spark a storage revolution.**

NAND Flash vs. Disk Drive

To see why NAND flash has such potential as a storage media for enterprise environments, compare its characteristics to those of a Fibre Channel disk drive. (See *Figure 1 – NAND Flash vs. Disk Drive*, on the next page.) The values shown are representative samples intended for making a ballpark comparison of the two technologies. While there are other types of drives and even other sizes and speeds of Fibre Channel with different characteristics, we chose a 146 GB, 15 RPM Fibre Channel drive because it offers high performance and is commonly used to store primary enterprise data, which is where NAND flash most obviously will compete, at first.

A single NAND flash memory stores 2.2 GBs compared to 146.8 GBs for a Fibre Channel drive, so another column was added to show values for the quantity of flash memories (67) needed to achieve an equivalent capacity. Since NAND flash are microchips, this quantity can be mounted on a circuit board and not differ significantly in size from a 3.5" disk drive.

Moreover, both NAND flash and disk drives offer persistent memory, which is the ability to retain data even when the power is off. This sets them apart from other forms of memory, like DRAM, that needs a continuous supply of power to retain stored data. With this fundamental similarity in mind, now consider the differences.

NAND Flash Strengths

First look at performance. The access time for a random read or write on a NAND flash is 220 times faster than the disk drive. **This is two orders of magnitude.** A disk drive is a physical device with rotating platters and a moving head, like a sophisticated version of a record player. In the computer age of silicon and electrons, it is amazing that a mechanical device has evolved and stayed with us so long, but the reality is that access performance is gated by the time it takes to spin the platter and move the head to the right spot. In contrast, NAND flash is 100% electronic and its access speed is determined by the flow of electrons through the circuit. A mechanical device simply cannot compete with an electronic one for random access performance.

The sustained transfer rate is a little closer, but NAND flash still comes out on top. A single flash memory is about a third as fast as the drive, but at equivalent capacity NAND flash is 22 times faster. This is an order of magnitude. Such a

IN THIS ISSUE

➤ NAND Flash vs. Disk Drive.....	1
➤ Fusion ioDrive	2
➤ Conclusion	3

Figure 1 — NAND Flash vs. Disk Drive
(Compare Shaded Columns)

CHARACTERISTIC	NAND Flash Memory (Single Unit)	NAND Flash Memory (Capacity Equivalent to FC Drive, at right)	Fibre Channel Disk Drive (15K RPM)
Quantity	1	67	1
Capacity	2.2 GB	147.4 GB	146.8 GB
Access Time (random read/write)	25 µsec	25 µsec	5.5 msec
Sustained Transfer Rate (sequential read)	33 MB/s	2,233 MB/s	100 MB/s
Operating Power (typical)	99 mW	6.6 W	16.5 W
Standby/Idle Power	132 µW	8.8 mW	10.7 W
Representative Price	\$30	\$2,010	\$325

Source: Hynix and Seagate

comparison implies the flash memories can work in parallel. As we will see later, this is possible and Fusion-io does it.

Power is another area where NAND flash shines. At equivalent capacity, operating power is 40% of the disk drive, which is significantly less, though on the same order of magnitude. It is idle or standby power where you see the big difference. NAND flash consumes 1,216 times less electricity during idle moments. **This is three orders of magnitude.** Disk drives keep spinning when idle, their mechanical motors consuming energy and generating heat. Since storage devices are not continuously accessed, if only during idle periods between I/Os, like stop-and-go traffic on busy city streets, the low idle power of NAND flash greatly reduces its overall power consumption. At a time when green computing and reduced energy consumption receives top billing for enterprise data center requirements, NAND flash has a clear advantage.

NAND Flash Challenges

The price of NAND flash per unit capacity is costlier than Fibre Channel drives by a factor of six. This difference in price could be considered worth it because of its performance and power advantages. In the coming years, NAND flash and disk drives may reach price parity because flash is riding a faster technology evolution curve. Moore's Law (or the spirit thereof) implies a continuity of semiconductor process and design improvements that are driving down the cost of flash – at a faster rate than costs are declining for disk drives.

Another challenge for NAND flash is the rated tolerance for program/erase cycles, which for this particular chip is 100,000. The rating means that bit errors are possible after a certain, albeit high,

number of writes. Therefore, controller logic must be able to intelligently detect and manage bit errors using techniques like error correction and wear-leveling¹. This same issue exists in disk drives and is manageable.

Fusion ioDrive

The Fusion *ioDrive* is a high-performance NAND flash storage system with the form factor of a PCI-e card. It provides access performance comparable to DRAM memory at capacities comparable to disk drives – which is what makes *ioDrive* a groundbreaking product.

ioDrive literally fits into a PCI-Express x4 slot of an enterprise server. The card accesses the server's internal bus directly for extremely fast, low-latency data transfer, as opposed to connecting via a NIC (network interface card) or HBA (host bus adapter). *ioDrive* is built with Fusion-io's proprietary architecture for high-density NAND flash called *ioMemory*. It uses a switched network protocol and 160 parallel pipelines to aggregate the performance of multiple flash memories. *ioMemory* also incorporates controller logic that intelligently manages bit errors in NAND flash, so the life expectancy of *ioDrive* is 8 years, versus 5 years for a disk drive, which is often reduced to a useful life of 3 years, due to the high out-of-warranty maintenance costs.

Capacity ranges from 80 to 640 GB. (See *Figure 2: Fusion ioDrive* below). A single card delivers 800 MB/s of sustained random reads, 600 MB/s of sustained random writes, and 100,000 IOPS. This level of performance is blistering and

¹ Distributing writes evenly, so blocks are not overused and subject to early degradation.

Figure 2 — Fusion ioDrive

CHARACTERISTIC	Fusion ioDrive
Capacities	80, 160, 320, or 640 GB
Sustained Read Transfer Rate	800 MB/s
Sustained Write Transfer Rate	600 MB/s
Sustained IOPS (8k random)	100,000
Operating Power (typical)	6 W
Approximate Price per GB	\$30

Source: Fusion-io

on par with a large SAN storage array, though the package fits on a PCI-e card! Multiple cards scale to higher capacity and performance or can be configured redundantly in a RAID configuration (no additional RAID card is required).

Fusion-io claims its fast speed can have a major impact on server and application performance:

- Serve 10 times the Web content
- Host 4 times the virtual machines
- Increase database transactions 25 times
- For simulations and visualization, accelerate data loading, saving and streaming 100 times

While actual performance gains depend on many system-specific factors, **these figures are at least within the realm of possibility for a key reason – computer systems have become increasingly bound by storage performance.** In spite of faster processors and buses and larger memory caches, servers frequently not reached their performance potential often because of storage access limitations. Remove that bottleneck and – *vroom* – the whole system takes off.

Another benefit is the nominal power consumption of ioDrive is 6 W, which is 40% of a Fibre Channel drive and an order of magnitude or two less than a SAN storage array. For data centers looking to minimize energy consumption, utility bills and heat output, ioDrive presents itself as a green alternative to disk. It is faster and much greener.

Since ioDrive is based on silicon rather than mechanical drives, it can withstand more vibration and shock, which is helpful for mobile environments and disaster resistance.

The target price range for ioDrive is approximately \$30/GB. A 320 GB card costs \$9,600 and a 640 GB card is \$19,200. If you consider this in terms of cost per IOPS, the price is significantly less than a disk drive or disk array.

ioDrive will be available in the first quarter of 2008. It will initially support *Linux* and then *Windows* servers, *XP*, and *Vista* three months later.

Conclusion

Fusion ioDrive is sparking a storage revolution that will shift the spectrum of tiered storage². Many enterprises now employ two types of disk arrays and perhaps even tape to optimize the cost of data storage. Primary data is stored on a high-performance storage tier based on Fibre Channel or SCSI drives, because enterprise applications need the speed. Secondary, backup, and archive data is stored on a lower-cost tier based on SATA drives or tape, because performance is less critical. Heretofore, solid-state disk based on SDRAM gained minimal traction at the top of the tiered storage spectrum because of its very high cost. Fusion-io's use of NAND flash in ioDrive to deliver a storage system with massive performance gains over disk at a much lower price than SDRAM, not to mention less costly than disk on a per-IOPS basis, suggest a shift is about to occur. **NAND flash-based enterprise storage systems like ioDrive will complement and even displace disk drives at the high end of the tiered storage spectrum. As NAND flash prices fall, this trend will only accelerate.**

The bottom line is that a storage revolution is underway. If you want to join this movement and enjoy the early gains, consider Fusion ioDrive.



² See the issue of *Clipper Notes* dated March 1, 2007, entitled *Tiered Storage Saves Money – Getting The Most Out of Your Storage Infrastructure* and available at <http://www.clipper.com/research/TCG2007032.pdf>.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "clipper.com" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "Navigating Information Technology Horizons", and "teraproductivity" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.