



High-Performance Storage for Hot Computing — IBM Introduces FlashSystem for Your Hot Data

Analyst: David Reine

Management Summary

Time is of the essence. It always has been that way. However, the tempo of life continues to accelerate. We used to be content to have our checks clear in 5 days and, now, we all are looking for near instantaneous deposits and payments. We used to be content getting two “turns” (program runs for debugging) per day and now we demand inline error correction and debugging. We used to be content to buy or sell a stock and have it clear in a week and now we expect it to happen in seconds. We used to be content to analyze sales data weeks and months after they took place and, well, now we want to be interactively suggestive during the sales process and immediately reflective thereafter. While it may be hard to project where the increase in tempo will take us in five-to-ten years, it is clear that there is little time in which to get things done.

Thus, it should not be a surprise that the need to “grease the skids” and move faster and in much higher volumes is one of the great IT challenges of the present. What was formerly “Good Enough” now may be a recipe for slowly going out of business. It seems as if nothing is good enough (for very long) in today’s business world. While this problem is not new, the magnitude and extent of it is without previous comparison. Going very fast is the new norm and tomorrow we will need to go even faster. What is an IT organization to do to satisfy those seemingly insatiable mission-critical applications that seem to be too hot to handle?

In the olden days (whether a couple of years ago or a couple of decades ago), the IT decision-makers would have held up their hand and said “No! What you want is just too expensive!” And that would have been the end of it, since the cost side of the equation drove the decision-making. Today, that answer is unacceptable, because the cost lost opportunities drives the equation. This does not mean that you should turn a blind eye to the scope of your costs, but that the costs must be weighed carefully with respect to your expected benefits, which may be plentiful.

When it comes to your most important applications, the demands for speedy performance tend to be at the highest level. Most of the time, if the application is mission critical, some if not all of the involved data is going to need to move very rapidly. It is this “extreme tempo” that puts the focus on IOPS (inputs/outputs per second). For these “hot” applications, time is money (with the emphasis on revenue and not on cost).

With the announcement of the *IBM FlashSystem Family*, IBM has developed flash optimized storage systems that strike an equal balance between performance and economics; that is, focusing on meeting the requirements for an extreme tempo, while also keeping an eye on the total cost of ownership. To learn more about IBM FlashSystem, please read on.

IN THIS ISSUE

➤ Conquering Data Access in the Data Center.....	2
➤ Flash to the Rescue.....	2
➤ IBM FlashSystem Family	4
➤ Conclusion.....	5

Conquering Data Access in the Data Center

Every data center, regardless of size, has been experiencing a period of unprecedented data growth, and the need for the immediate analysis of that data on a 24x7x365 basis is, almost, universal. It doesn't matter whether your enterprise needs to store ten terabytes of data or ten petabytes, in order to gain a competitive advantage, you need a scalable, reliable, efficient system that will enable your storage to scale to meet the growing storage needs of capacity, performance, and reliability created by the "Age of Big Data". Most data centers are doubling their storage requirements every 12-to-18 months, while some are experiencing even greater growth, compounding the cost, management complexity, and the issues involved in extracting valuable information from that data on a 24x7x365 basis. It is critical – mission critical – for the data center to be able to identify the "hottest" data and move it front-and-center to the highest performing storage. This is where the enterprise data center needs some new technology to accelerate their data processing.

Accelerating data retrieval time from milliseconds to microseconds can mean the difference between seconds and minutes for applications, such as real-time analytics and mission-critical workloads such as OLTP, ERP, order processing, cloud inquiries and financial transactions. Assigning the hottest data to the fastest storage is a critical concern for the IT staff. How can the IT staff separate the wheat (hot data) from the chaff (cold data) and where do they put that wheat?

Today, many enterprises literally have thousands of users trying to access critical, shared data from multiple applications. Because of the need for immediate access to that data, the IT staff has to distribute it over significantly more hard disk devices (HDDs) than should be required to store it, in order to ensure the fastest response. This dramatically increases the total cost of ownership (TCO) for the primary storage (on a per-GB basis). The data center must improve the increased access to storage while at the same time maintaining SLAs for availability, data integrity, security, and data flexibility. Deploying five-to-ten times (or more) the amount of capacity wastes resources, both in terms of capital expenditures and operational expenses, such as floor space and energy consumption. It behooves the IT staff to reduce, or

eliminate, the need to waste up to 90% of any deployed asset, in this case storage configured to minimize the mechanical actions of disk drives, both rotational movement and the positioning of the read/write heads. Many enterprises today also have active programs to make their data centers "greener". This can be advanced by reducing the unused infrastructure and deploying more energy efficient solutions in the data center.

In addition to the implementation of a more energy efficient infrastructure, it is mandatory that mission-critical data be deployed on the most reliable and durable storage, in order to ensure the availability of that data to enable access to the business insights that can be learned from the volumes of data being collected. This requirement demands an environment with no single point of failure. A fast, efficient, reliable infrastructure leads to a smarter enterprise that can be early to market.

How does the enterprise calculate the TCO of the storage infrastructure? Is determining the cost of storage per MB an adequate starting place? What about the cost of storage per I/Os per second (IOPS)? By improving the utilization of the storage being deployed, the IT staff can, in fact, reduce the TCO, saving the enterprise Big \$\$ for their Big Data. Remember, even if the data center has a higher acquisition price for storage, there may be a lower overall TCO.

One way to stop, and perhaps reverse, the ever-increasing TCO of your storage infrastructure is to change your existing storage paradigm and transition to a hybrid solution with both HDDs and flash or other solid state devices (SSDs) comprising the storage configuration.

Flash to the Rescue

Can flash storage (SSDs or other flash devices) provide mission-critical applications with the performance, reliability, availability, and security they need without literally destroying the IT infrastructure budget? The answer to that is an unequivocal "YES"; however, the IT staff must be careful in the selection of flash infrastructure.

We all may have been created equal; but quite clearly, this is not true for solid state storage technology. While it is true that all solid state storage platforms have been designed for speed, some also have been designed to provide enterprise-class SLAs, as well. Consumer-grade

SSDs, such as those designed for laptops and other portable devices, must be power-efficient, as well as high performant, in order to preserve battery life. Enterprise-grade flash technology, however, must have the qualities enumerated above, in order to guarantee continuous access to that mission-critical data. Unfortunately, there is a price to pay for that access and reliability. Thus, it shouldn't be a surprise that enterprise-grade flash technology is more expensive than consumer-grade SSDs!

The IT staff needs to determine the correct mix of HDDs and flash that is appropriate for their application mix. Quite clearly, enterprise-grade flash has the qualities necessary to respond to the key requirements of your most mission-critical applications need for hot data at a high tempo.

First of all, if your applications are not I/O bound, then HDDs may be able to provide an adequate performance level at a lower cost than flash. While capacity per drive, storage density, and cost/GB have all steadily improved over time, access time to your critical data has not, because you are trying to do so much more and usually in less time than before. Because of the physical limitations of HDDs, specifically, rotation speed and movable heads, improvement in response time and overall performance per drive essentially has been stagnant. With ever-improving CPU performance and memory speed solving part of the tempo problem, HDDs now are creating bottlenecks to enterprise success.

Traditionally, in order to achieve the desired number of IOPS required to meet enterprise SLAs, data has been distributed over many more devices than required in order to put more read/write heads into play. In fact, many mission-critical HDDs have been configured to use only 10% to 20% of their native capacity in order to improve response time. While HDDs may carry a relatively low cost/GB, making them an ideal target for your high-capacity "cold" data, they carry a very high cost/IOPS for your urgently required "hot" data.

Flash technology, on the other hand, significantly can improve performance for that "hot" data in I/O bound applications, without application changes, through better latency and significantly improved IOPS. Depending upon the application and configuration, SSDs and flash-optimized systems can provide hundreds of times more in terms of IOPS, working at speeds closer to that of server memory. While SSD

cost/GB generally carries a higher cost per device, reducing the number of HDDs required by as much as 10:1 usually will more than compensate on a TCO basis. Cost/IOPS, however, is always significantly lower for SSDs and other flash technology.

A hybrid environment with both HDDs and flash storage provides the ideal configuration for application access to both cold and hot data. With flash storage deployed, the improved density through better storage utilization reduces the number of drives required to achieve your performance target, thus also reducing the amount of floor space and energy required. Thus, deploying flash storage can actually lower the TCO of the storage infrastructure.

As mentioned above, there are two classes of SSDs, possessing functionality differences similar to the differences in HDD functionality¹. These are referred to as consumer grade and enterprise grade SSDs. Consumer-grade SSDs provide relatively high performance with, significantly reduced energy consumption, in order to extend the battery life of mobile devices. Enterprise-level SSDs also possess these features, as well as dual porting, end-to-end data integrity to protect data on the fly, and rigorous testing of all components to ensure interface compatibility, and ensuring data availability across multiple devices. In enterprise environments, SSDs may be appropriate as a first step to build hybrid storage infrastructures – but note that there are now more integrated solutions available that maximize the benefits of flash devices.

For several years, flash SSDs have been available in two technologies: *single-level cell (SLC)* and *multi-level cell (MLC)*.² SLC flash is very fast, with a longer life expectancy and higher reliability, but more cost than MLC, which has lower capacity and higher power consumption. Reliability issues largely can be overcome by the design structure of the SSD, such as interleaving, changes to writing algorithms, and over-provisioning (providing excess capacity) with which the wear-leveling algorithms can work. Also, by increasing error correction codes, you can compensate for the lower reliability and extend the life of MLC devices.

¹ The IT staff has always had different functionality available from different speed SATA, SAS, and FC HDDs.

² See the issue of *Clipper Notes* dated February 10, 2009, entitled *A New Tier of Storage Appears – Faster, Solid-State Drives State Their Case*, and available at <http://www.clipper.com/research/TCG2009006.pdf>.

MLC flash deployments in the enterprise generally now have transitioned to Enterprise MLC (eMLC), a type of MLC flash that has been enhanced to accommodate more write cycles than consumer-grade MLC flash drives. Some variants of eMLC flash are rated at 30,000 write/erase cycles, compared to 1000 to 3000 write/erase cycles for consumer-grade MLC. The problem is that flash memory wears out, and can only sustain so many erase/program cycles before writes no longer can be guaranteed.³ Today's eMLC flash technology provides higher durability through more write/erase cycles than generally available from consumer-grade MLC, thus adding enterprise-class performance and reliability to the more cost-effective characteristics of MLC. While the cost of flash is going down, capacity and IOPS continue to rise. This enables the IT staff to reduce the size of the storage infrastructure by moving the hottest data to MLC drives and repurposing or eliminating the thinly-used high-performance drives where the data previously resided.

IBM recently introduced FlashSystem, a family of all-flash storage systems based upon technology acquired by IBM in the acquisition of Texas Memory Systems. FlashSystem provides turnkey access to the benefits of flash (at greater scale than individual SSDs) – to enable the data center to perform advanced data analytics faster than ever before – in order to gain a competitive advantage.

IBM FlashSystem Family

The IBM FlashSystem entails a family of affordable SSD flash storage devices with specific capacity and performance configurations to meet a variety of application requirements. IBM FlashSystem has an innovative design with no single point of failure, system and application integration, with management tools to enable effective use.

Based upon tests conducted by IBM at their Almaden Research Lab, IBM has determined that flash storage solutions deliver 6.7 times more capacity in a single rack, 19 times cost efficiency in \$/IOPS, and 115 times more energy efficiency than thinly-provisioned HDD solutions. IBM has eliminated storage bottlenecks

³ What this means is that once the maximum has been reached, you really shouldn't be writing data to that location any more.

with IBM *MicroLatency*⁴, delivering access times less than 100 microseconds, to enable faster decision making. This enables your mission-critical applications to operate with an access time much closer to that of CPUs, DRAM, and internal buses in the I/O data path.

Many applications, such as *DB2*, are integrated closely with IBM storage technologies, such as the *IBM System Storage SAN Volume Controller (SVC)*⁵ and *IBM System Storage Easy Tier*⁶, to manage data placement between faster and slower tiers of storage devices – for example, FlashSystem and HDDs. IBM has estimated for some workloads that with only 10% of the data deployed on flash, IBM's Smart Data Placement can achieve 90% of the throughput increase of an all-flash system, for maximum performance, at a significantly lower capital expense.

The IBM FlashSystem Family includes two systems configured with SLC flash and two with eMLC flash, each delivered in a 1U rack-mount format. The IBM FlashSystem 810 and FlashSystem 820 have been designed with eMLC flash, which is changing the economic paradigm of flash by delivering extreme performance with higher density, and a low cost per terabyte. In fact, eMLC flash chips offer five-to-ten times more endurance than consumer-grade MLC at the chip level.

On the other hand, the IBM FlashSystem 710 and FlashSystem 720 are deployed with SLC flash, providing the highest performance and lowest latency for enterprises looking for the ultimate in performance. All FlashSystem products have a maximum bandwidth for reads of 3.3 GBs/second for FC and 5 GBs/second for InfiniBand.

IBM has designed the FlashSystem to reduce business interruptions, without sacrificing performance or usable capacity, via *IBM Variable Stripe RAID*⁷ technology to help provide a

⁴ MicroLatency enables faster decision making by facilitating an extreme-performance data path to accelerate critical applications and help users achieve a competitive advantage

⁵ See [The Clipper Group Navigator](#) entitled *IBM Upgrades SVC with Solid State Drives – Achieves Better Storage Utilization* dated October 19, 2009, and available at <http://www.clipper.com/research/TCG2009046.pdf>.

⁶ See [The Clipper Group Navigator](#) entitled *Maximizing the Business Value of Information and Lowering Energy Consumption with IBM's DS8700* dated April 13, 2010, and available at <http://www.clipper.com/research/TCG2010018.pdf>.

⁷ VariableStripe RAID, a technology obtained with the acquisition of Texas Memory Systems, enables continued operation

higher level of protection from flash device failure. In addition, FlashSystem utilizes standard chip-level RAID technology, enhanced error-correcting code and internal redundancies to minimize the requirement for maintenance, freeing the IT staff for more strategic activities, such as accelerating decision making and implementing I.T. efficiency, achieving a high bandwidth and an ultimate IOPS performance level. Additional benefits occur from cost efficiencies gained through reduced power consumption and floor space.

FlashSystem 820

The FlashSystem 820 is a scalable storage device, providing up to 24 TBs usable eMLC flash storage (up to 33 TBs raw). It has a minimum write latency of 25µs and 110µs for reads. With 4KB blocks, the maximum IOPS in a pure read environment is 525,000, with 280,000 IOPS in a pure write environment, all while consuming only 300 watts. Its maximum bandwidth for writes is 2.8 GBs/second for both FC and InfiniBand. The FlashSystem 820 can be configured for either RAID 0 or RAID 5

FlashSystem 810

The FlashSystem 810 is a scalable storage device, providing up to 10 TB usable eMLC flash storage (up to 13.7 TB raw). It has a minimum write latency of 60µs and 110µs for reads. With 4KB blocks, the maximum IOPS in a pure read environment is 550,000, with 400,000 IOPS in a pure write environment, all while consuming only 350 watts. Its maximum bandwidth for writes is 2.8 GBs/second for both FC and InfiniBand.

FlashSystem 720

The FlashSystem 720 is a scalable storage device, providing up to 12 TBs usable SLC flash storage (up to 16.5 TBs raw). It has a minimum write latency of 25µs and 100µs for reads. With 4KB blocks, the maximum IOPS in a pure read environment is 525,000, with 400,000 IOPS in a pure write environment, all while consuming only 350 watts. Its maximum bandwidth for writes is 3.3 GBs/second for both FC and InfiniBand. The FlashSystem 720 can be configured for either RAID 0 or RAID 5.

FlashSystem 710

The FlashSystem 710 is a scalable storage device, providing up to 5 TBs usable SLC flash

storage (up to 6.9 TBs raw). It has a minimum write latency of 60µs and 100µs for reads. With 4KB blocks, the maximum IOPS in a pure read environment is 570,000, with 400,000 IOPS in a pure write environment, all while consuming only 280 watts. Its maximum bandwidth for writes is 3.3 GBs/second for FC and 4.5 GBs/second for InfiniBand.

Obviously, the determination of which of these systems to deploy will depend upon the application requirements for throughput and reliability and tolerance for acquisition cost.

Conclusion

If your enterprise is looking for an immediate return on its investment for accelerated storage while improving its energy efficiency and lowering the TCO of the IT infrastructure, then the move to flash technology with solid state storage may make sense.

IBM's FlashSystems can be used to reduce/eliminate business interruptions caused by insufficient IOPS without sacrificing performance, reliability, or capacity. They enable the enterprise IT staff to accelerate decision making across all mission-critical applications with improved IT efficiency and a high bandwidth to assist the data center in improving online analytical processing by speeding access to hot data.

If these issues sound too familiar to you, then take a look at IBM's FlashSystem technology. It may be the answer you have been seeking.



when a flash plane fails allowing for longer mean time between failures.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-first year, is an independent publishing and 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

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- **Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” 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, The Clipper Group Calculator, 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.

Disclosures

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper’s web pages, posting of Clipper documents on other’s websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts’ travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper’s analysts.

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.