An Assessment of SSD Performance in the IBM System Storage DS8000

June 5, 2009

Executive Summary

Solid State Drives (SSDs) enable dramatically higher throughput and lower response times, providing the potential to significantly lower operational costs in the data center. IBM recently announced the IBM System Storage DS8000 Turbo series with Solid State Drives (SSDs).

This paper will discuss the advantages of SSD, highlight SSD best practices, provide an assessment of performance using DS8000s with SSD, and demonstrate energy, cooling, and space savings with SSD.

Why SSD

Solid State Drives (SSDs) are increasingly becoming a very attractive option for enterprise storage needs. SSDs have no moving parts so they perform at electronic speeds without the mechanical delays associated with traditional spinning Hard Disk Drives (HDDs). Because SSDs enable dramatically higher throughput and lower response times than HDDs for random I/O, they provide the potential to significantly lower operational costs in the data center despite higher current acquisition cost per GB. To realize these benefits, it is critical to specifically target usage to applications that require high IOPS/GB and/or low response times.

The IBM DS8000 now supports three performance tiers of storage:

- Tier 0: SSDs. Highest performance and cost/GB
- Tier 1: 15K RPM HDDs. High performance and lower cost/GB
- Tier 2: 7200 RPM HDDs. Lowest performance and cost/GB

Previously, IT managers needed to obtain large quantities of 15K RPM HDDs for certain applications and use only a small portion of the capacity of each HDD to meet their performance requirements. This practice can be costly as it reduces capacity utilization. Now a large number of HDDs may be replaced with a small number of SSDs, fully utilizing the capacity of each SSD and realizing improved system performance while also saving on space, energy, and cooling.

The next sections of the paper will discuss SSD best practices and the performance that can be achieved with SSD.

SSD Best Practices

Using the right mix of storage drives will provide optimal performance at the minimum cost, energy, cooling and space usage. To maximize the benefit of SSDs it is important to place only data which requires a high throughput and low response time on them. This data is referred to as "hot" data. Data that requires lower throughput is referred to as "cold" data. Once "hot" data is moved to SSD, the remaining data may be "cold" enough to allow moving a large portion of it to high capacity 7200 RPM HDDs and still meet the required performance.

Determining the "temperature" of data and moving it to the proper tier can be difficult. Performance management tools across platforms may be useful in assisting IT managers with these decisions.

AIX on IBM Power Systems provides performance tools that can be used to determine if a configuration has hot data that would perform better if moved to SSDs. The IBM System z I/O architecture provides a detailed breakdown of time spent executing I/O operations. SSDs are ideally suited to benefit workloads that are incurring high numbers of cache misses (for example, random reads), which can be determined from this data.

The DS8000 also provides the ability to obtain cache statistics for every volume in the storage system. These measurements include the count of the number of operations from cache to the backend storage, the number of random operations, the number of sequential reads and sequential writes, the time to execute those operations, and the number of bytes transferred. Volumes with the highest device loads may be considered for full volume migration to SSD. New z/OS tooling can identify the hottest data sets on the most stressed volumes. Movement of individual data sets from these volumes could be considered as an alternative to full volume migration when SSD space is limited.

In order to achieve maximum performance from SSDs, the input from performance tools and analysis can be used along with the SSD best practices outlined below.

- Place "hot" data on SSDs, "warm" data on 15K RPM HDDs, and "cold" data on 7200 RPM SATA HDDs.
- Use SSDs for applications that require low response times and are cache unfriendly.
- Place random data on SSDs and sequential data on HDDs.
- Use SSDs for applications that traditionally "short stroke" (use a small portion of the capacity of) large numbers of 15K RPM HDDs with low capacity utilizations.
- Consider using a smaller storage cache when using SSDs than might be used when using HDDs. For hybrid DS8000s containing both SSDs and HDDs, it is advisable to use the same size storage cache as for a configuration of all HDDs so that read hit ratios on the volumes placed on the HDDs are not reduced.
- Perform the appropriate capacity planning before placing SSDs into a remote copy environment. If SSDs are used for remote copy source volumes they should also be used for the remote copy destination volumes. If not, then the secondary HDD based destinations may become the bottleneck in the system.
- Use SSDs with FlashCopy either for source or destination volumes. If SSDs are used for source volumes while HDDs are used for the secondary, it is a good idea to do the FlashCopy with background copy and during a period when the write rate to source volumes does not exceed the capability of the destinations.
- Use High Performance FICON (zHPF) with SSDs for higher throughput and additional reduction in the total response time.

Performance Results for SSD

Performance results in this section compare SSDs with HDDs in the IBM DS8000. IBM System z measurements were taken on the DS8000 with SSDs and HDDs to compare performance.

The results in Figure 1 were measured by a DB2 I/O benchmark. They show random 4KB read throughput and response times. SSD response times are very low across the curve. They are lower than the minimum HDD response time for all data points.

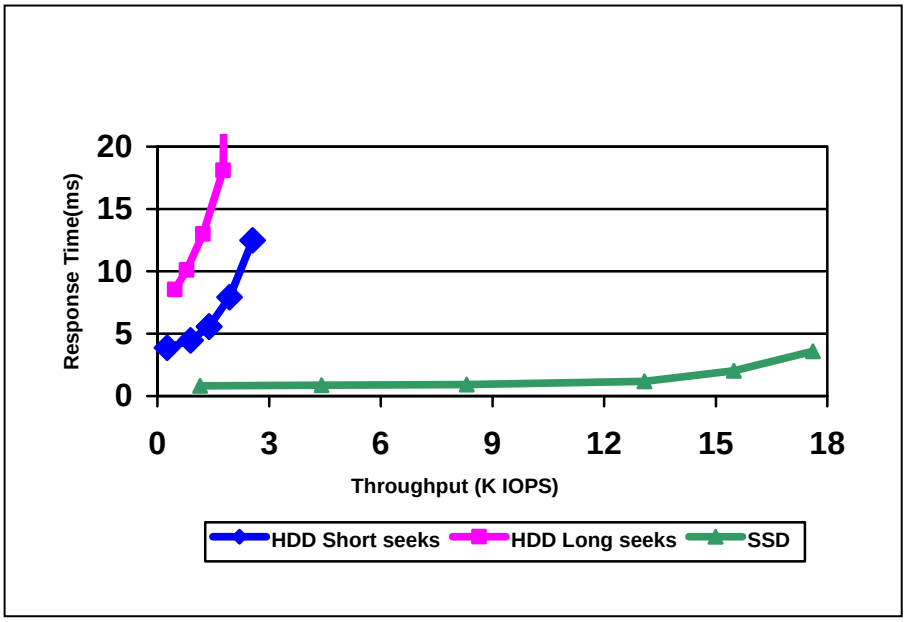

Figure 1 - DB2 on CKD Random Read Throughput/Response Time Curve

Figure 2 shows that SSDs provide approximately the same improvement on random writes as they do on random reads. Note that random write performance is lower than random read performance on HDD and SSD due to the extra drive I/Os done on RAID5 writes.

Figure 2 - FBA 4KB Random IO: SSD vs HDD on one RAID5 Rank

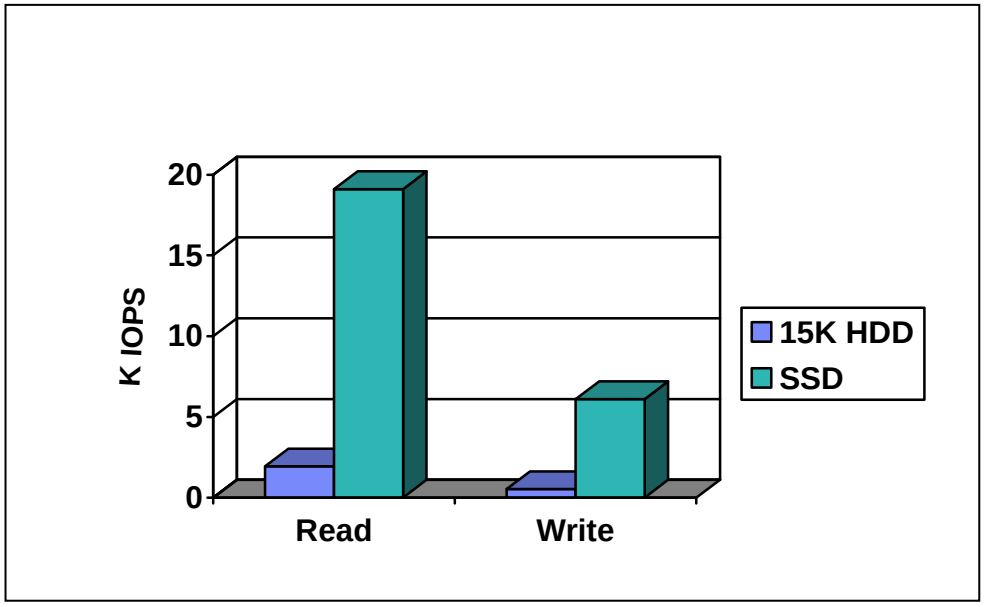

For random I/O, SSDs provide much higher throughput at a much lower response time. The SSDs supported in the DS8000 are so fast that the device adapter may become the performance bottleneck on some random workloads. For sequential I/O, the device adapter was already the performance bottleneck with HDDs so using SSDs for sequential I/O would not provide a substantial increase in performance.

Also note that:

- Applications that require low response times may not be able to meet their requirement with HDDs no matter how much they short stroke their HDDs. SSDs are a good match for these applications.

- Applications that traditionally use a very large storage cache may now be able to use a combination of SSDs and a small cache and save on the capital cost and energy usage of the large cache.

- zHPF has a lower response time than standard FICON. When performing either cache hits or I/O to SSDs, using zHPF provides a significant additional reduction in the total response time.

More Advantages of SSD

In addition to the dramatic performance advantages SSDs provide over their HDD counterparts for transaction-intensive applications, SSDs boast other key advantages, such as reliability, lower energy usage, less cooling requirements, and the ability to reduce data center footprints. When combined, these advantages can add up to significant performance improvements as well as a lower costs structure for business critical applications.

Long term, SSD has the potential to become much more reliable than HDD since there are no moving parts. IBM is working with industry leading device providers to identify and leverage the best technology to provide performance, reliability, availability, serviceability and other attributes for leadership devices in IBM systems.

Since deploying SSDs can eliminate the expensive habit of "short stroking" HDDs to enable higher throughput performance for critical applications, clients may see a considerable reduction in their storage footprints. Recall that "short-stroking" HDDs require clients to use a small portion of the HDDs capacity, which is the tradeoff for higher performance. By eliminating this tradeoff, virtually 100% of the SSD is utilized, which can greatly reduce the number of drives needed. As clients continue to struggle with managing the tremendous growth of data to manage, more efficient storage utilization may pay big dividends, especially in metropolitan areas where real estate values are at a premium.

With respect to energy, each SSD uses approximately half of the energy of a 3.5" 15K RPM HDD. For applications that are able to replace large numbers of HDDs with a small number of SSDs, energy savings are compounded.

The technical brief "*IBM System z® and System Storage DS8000: Accelerating the SAP® Deposits Management Workload With Solid State Drives*" provides an example of the potential energy, cooling and space savings from SSDs. In this study, the hybrid SSD/HDD configuration provided the following benefits:

- 22% higher throughput at 50% lower response time
- 60% floor space savings
- 74% electrical power and cooling savings

In fact, using the SSD/HDD configuration, approximately 22.9 kilowatts were saved. At 15.78 cents per kilowatthour (New York, all sectors, 1/2009), savings resulted in approximately $31,623 per year.

Conclusion

SSDs are an emerging technology for enterprise storage clients that can show immediate benefits in terms of performance as well as other operational characteristics**.** Given the distinct attributes and costs of SSDs and HDDs, it is clear that both drive types will coexist for some time. This coexistence will require a strong focus on smart data placement and the subsequent data migration, which are two strategic areas for IBM SSD solutions.

SSDs have no moving parts and provide much higher throughput and much lower response times for random I/O than traditional spinning HDDs. They can also significantly lower operational costs in the data center. Since SSDs currently have a higher cost per GB than HDDs, they are specifically targeted at applications that require high IOPS/GB and/or low response times and may eliminate the practice of "short stroking" for these performance critical applications.

By eliminating the seek times of their spinning counterparts and providing direct access to data, SSDs may dramatically boost performance and allow clients to maximize drive capacity utilization. This consolidation may enable the replacement of a large number of HDDs with a much smaller quantity of SSDs and can also lower cache memory requirements while maintaining the best levels of response times. By reducing energy consumption, cooling expenses and floor space costs, SSDs are an important part of the future in enterprise storage.

References

System z® and System Storage DS8000: Accelerating the SAP® Deposits Management Workload With Solid State Drives
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101442

Stop spinning your storage wheels: z/OS Support for solid state drives in the DS8000 storage subsystem in z/OS Hot Topics Newsletter, Issue 20
http://www.ibm.com/systems/z/os/zos/bkserv/hot_topics.html

IBM Corp. 2009. .IBM System Storage DS8000 Turbo series
ftp://ftp.software.ibm.com/common/ssi/pm/sp/n/tsd00374usen/TSD00374USEN.PDF

IBM Corp. 2009. US Announcement Letter 109-120: IBM System Storage DS8000 series (Machine types 2421, 2422, 2423, and 2424) delivers new security, scalability, and business continuity capabilities.
http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS109-120&appname=USN

SSD Performance and Power Advantage whitepaper, STEC Inc.
http://www.stec-inc.com/green/green_ssdsavings.php

Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State:
http://www.eia.doe.gov/cneaf/electricity/epm/table5_6_a.html

IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
June 2009

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided “AS IS” and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

When referring to storage capacity, 1 TB equals total GB divided by 1000; accessible capacity may be less.

The IBM home page on the Internet can be found at: http://www.ibm.com.

The IBM Power Systems home page on the Internet can be found at: http://www.ibm.com/systems/power/

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of the Microsoft Corporation.

Intel, Itanium and Xeon are registered trademarks and MMX and Pentium are trademarks of Intel Corporation in the United States and/or other countries.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. In the United States and/or other countries.