Disk-application performance gap and the future of SSS

Silverton Consulting, Inc.
StorInt™ Briefing

Semiconductor storage systems (SSS) have recently emerged to challenge disk subsystems and are enjoying moderate success in very high performance environments. Despite disk's current dominant position in the storage market, Flash SSS storage vendors such as Texas Memory Systems (TMS) have a unique opportunity today to break out of this high performance niche to address more normal application requirements. In fact, a gap has already begun to develop between the storage performance requirements of common applications such as MS Exchange Server and the ability of traditional hard disk systems to meet these performance requirements. While data centers have been bandaging this problem by over-provisioning their disk storage, at some near point Flash SSS storage will become decisively competitive.

For an analogy to the application performance gap, consider automobile miles per gallon (MPG) vs. miles per tank. Although automakers can easily increase miles per tank (by increasing tank capacity), improving MPG takes much more effort and has a much more direct impact on the cost of driving. Similarly, increasing hard drive capacity can cause drive performance to increase but does little to address increasing application performance requirements.

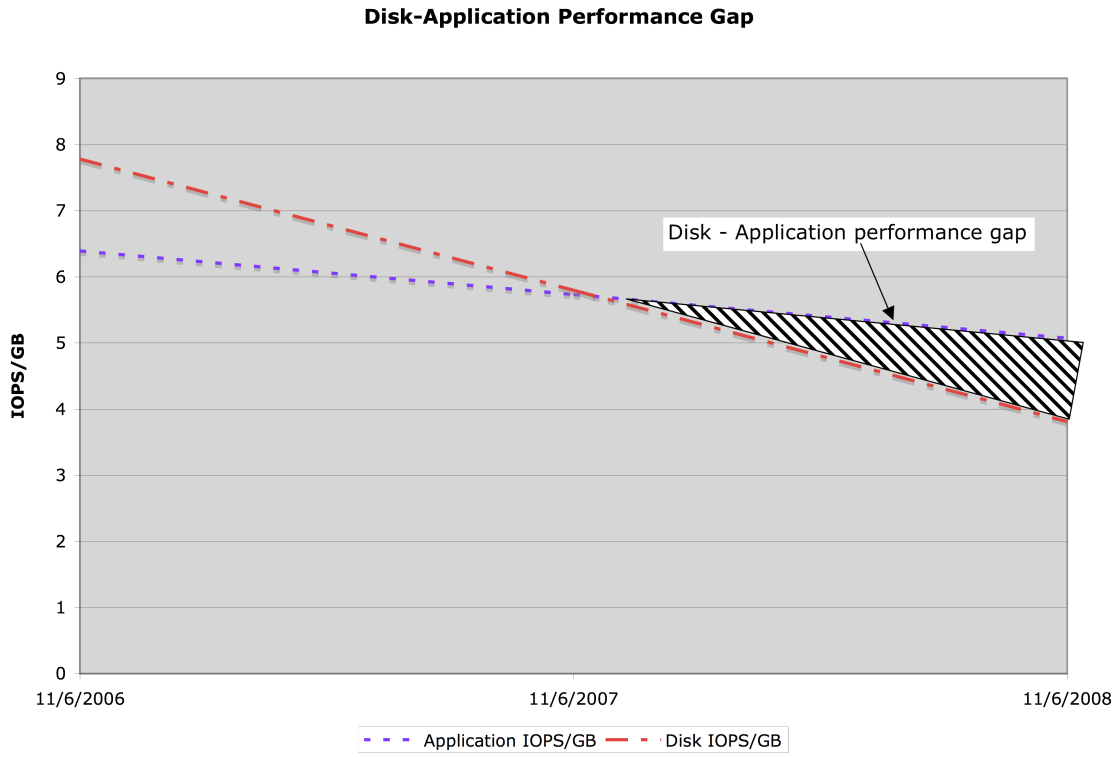

Figure 1 Disk-application performance chart

What's happening to disk performance?

An analysis of disk subsystem performance as reported by the Storage Performance Council (SPC), an internationally recognized and respected independent benchmarking organization, leads to some unexpected results. For example, analyzing SPC-1 benchmark results, designed to measure maximum sustainable IO operations per second, at least two surprising conclusions can be drawn. Specifically:

- Average disk subsystem I/O performance, while in aggregate continues to improve, is on a steady decline on a capacity basis, i.e. IO operations per second per GB (IOPS/GB), and is now running below 4 IOPS/GB of user storage.

- Average disk subsystem costs, while continuing to decrease on a subsystem cost per GB (SS$/GB) basis, are relatively flat on a subsystem cost per drive (SS$/Drv) basis and hovering at ~$2850 SS$/Drv.

The above disk storage trends are identifiable by plotting representative SPC-1 result histories as shown in figure 2[1]. For example, a plot using linear regression of SPC-1 results for IOPS/GB over time for enterprise disk subsystems[2] shows a rapid, disturbing decline in disk performance on a capacity (per GB) basis. Indeed, this analysis shows that disk subsystem performance has decreased 21% each year of reported SPC-1 results. Extrapolating to November 2008, the average SPC-1 disk subsystem performance would fall to ~3.8 IOPS/GB.

[1] Storage Performance Council results are available from http://www.storageperformance.org as of November 8, 2008

[2] Enterprise class subsystems are defined as SPC-1 results from disk subsystems utilizing more than 100 disk drives.

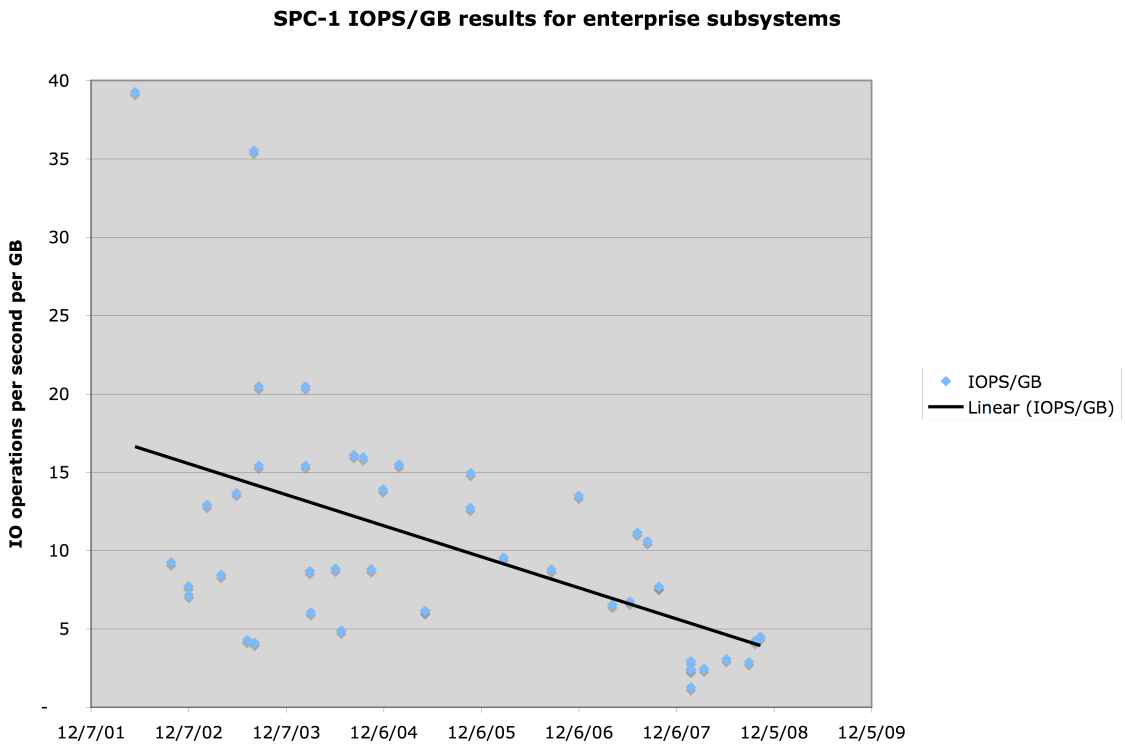

Figure 2 SPC-1 IO operations per second per GB

On the other hand, a plot of SPC-1 results charting IO operations per second per drive (IOPS/Drv) performance shows an upward trend (See figure 3). In fact, the analysis indicates a 5% boost to subsystem drive performance to ~204 IOPS/Drv for the SPC-1 enterprise class results.

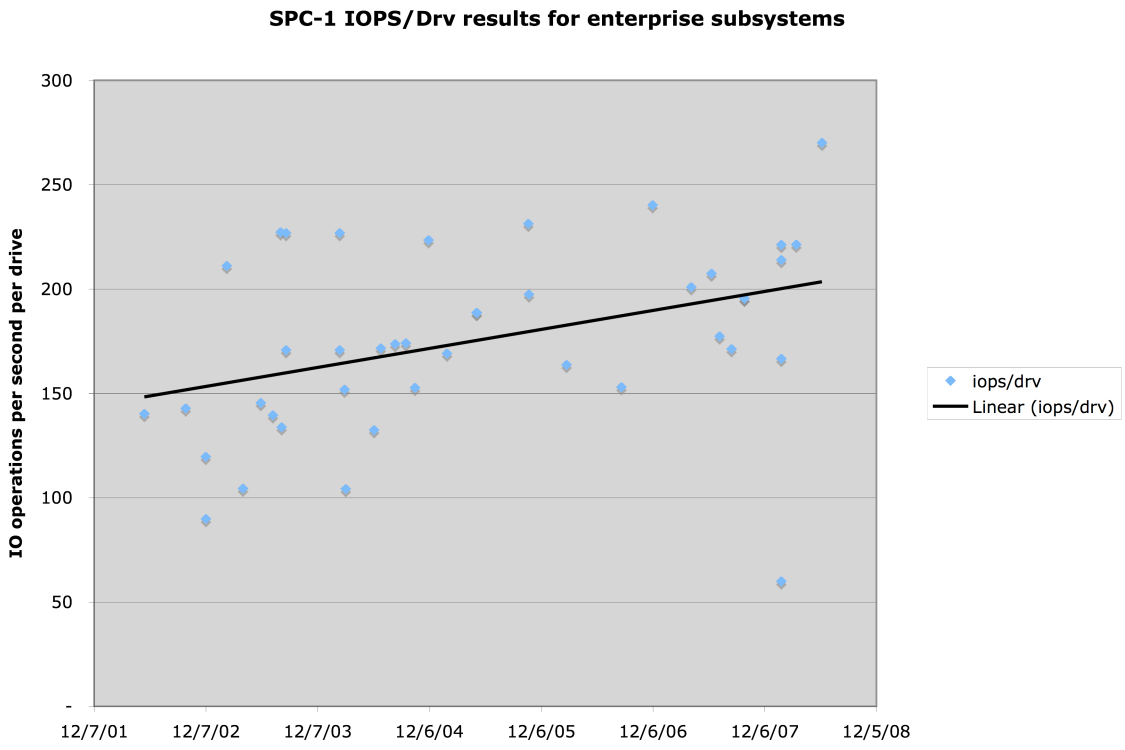

Figure 3 SPC-1 IO operations per second per drive

At first blush, the two trends illustrated in figures 2 and 3 might seem contradictory. In fact, they are not. Drive capacity has been almost doubling every two years over the last decade. This doubling in capacity has offset the 21% performance degradation at the purely IOPS/GB level and has allowed disk vendors to tout overall performance improvement.

What's happening to application performance?

Surprisingly, but at least under the Microsoft Exchange Solution Review Program[3] (ESRP), a linear regression analysis of the reported data indicates a 9% per year decrease over the last two years in database and log performance requirements. In fact, current ESRP results indicate a ~5.0 IOPS/GB. Mapping the 9% degradation of transfer requirements to the 21% downward trend of disk subsystem performance results in an ever-expanding application performance gap as illustrated in figure 1. The widening gap between disk performance and application requirements can most effectively be addressed by higher performance storage such as flash SSS.

[3] ESRP V2.1 results are available at http://technet.microsoft.com/en-us/exchange/bb412164.aspx as of 11/08/2008

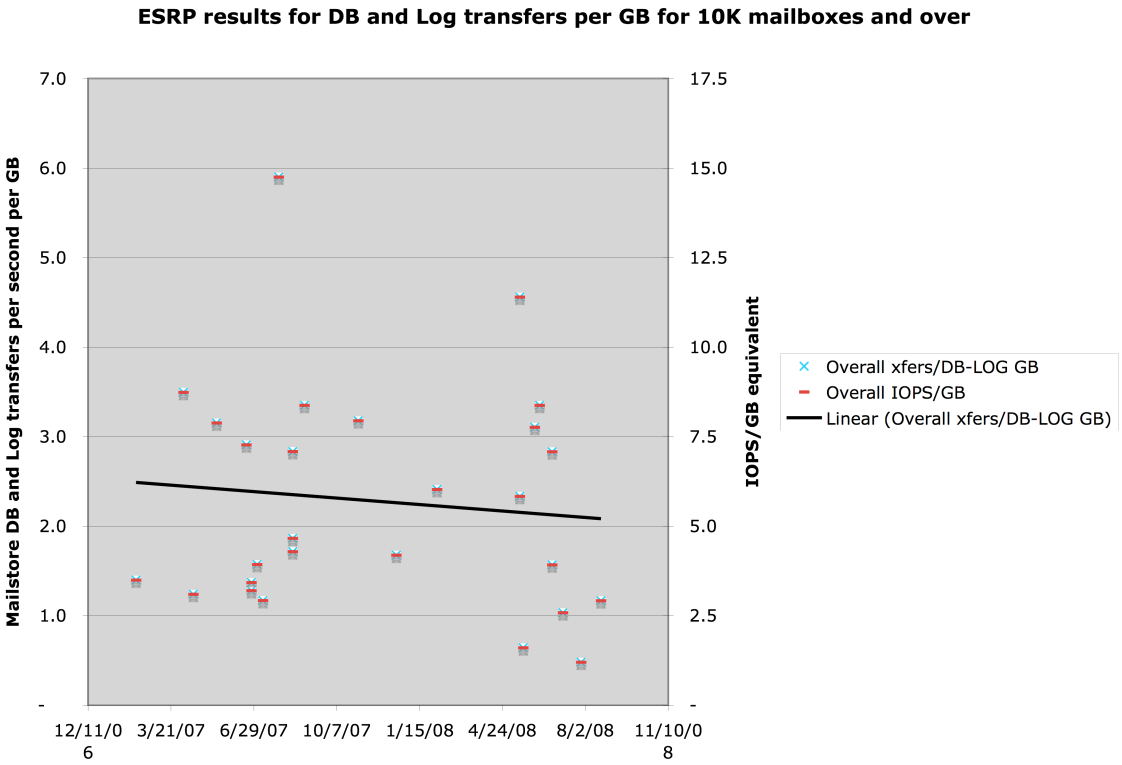

Figure 4 ESRP reported database and log transfers per GB

Arguably, the choice to use Microsoft's ESRP information may or may not be indicative of a particular application's performance requirements. However, ESRP information was chosen because:

- The expected universality of overall email applications
- The widely publicized availability of results
- The ability to extract results reporting only 10,000 or more e-mailboxes to equate to enterprise-like SPC-1 results
- The ability to compare results between ESRP and SPC-1 using a common factor based on IOPS/GB (based on a 2.5 multiplier to equate ESRP database and log transfers to IOPS/GB).

What's happening to disk storage costs?

Traditionally storage vendors have over-provisioned disk storage to increase subsystem performance. That is, an overage of disk drives allows I/O to be spread over more drives and thus supports more performance. While this practice may be expedient, it is never optimal and may in fact be uneconomical sooner than expected.

An analysis of SPC-1 enterprise disk subsystem costs indicates a ~2% annual rise (SS$/Drv) to a fully burdened $2,845 cost per drive. In contrast, by expanding drive capacity, vendors have decreased the SS$/GB ~24% annually to a current ~$48 fully burdened cost per GB.

Does SSS make sense yet?

Unfortunately, flash SSS performance and pricing is not as widely reported as disk storage information. In fact, flash SSS subsystem vendors have yet to post a single SPC-1 benchmark result. Thus, in comparing flash SSS to disk subsystems, heavy reliance must be placed on vendor supplied and not independently verified information.

Given this caveat, one flash vendor claimed 10,000 random read operations per second but only 250 random write operations per second with their flash drive across 64 GB of flash storage. At the raw device level, this performance equates to 156 reads/GB or 4.0 writes/GB. Using a widely accepted, historical read-to-write ratio (R:W) of 2:1, this flash device could easily sustain 8 reads plus 4 writes per second or a remarkable ~11.2 IOPS/GB.

While this performance is outstanding, a prospective buyer would still be paying a premium on a SS$/GB basis alone. For example, TMS currently sells their 2TB RamSan flash storage device for $150,000 or $75/GB. However, comparing $ to IOPS/GB, or a price-IOPS equivalent, SSS devices fare much more favorably.

For example, using TMS $75/GB cost against the current drive cost of $2845, a price equivalent of ~33GB of flash storage is derived. By dividing this $33/GB price equivalent into the ~204 IOPS/Drv for enterprise disk storage, a price equivalent breakpoint of 6.2 IOPS/GB is calculated.

As such, at least for flash SSS performance and TMS pricing information:

- If an application requires more than 6.2 IOPS/GB of performance it would be cheaper today to buy flash SSS storage to sustain this performance.
- If an application requires less than 6.2 IOPS/GB of performance it would be cheaper today to buy disk storage to sustain this performance.

While it is clear that 6.2 IOPS/GB represents superior performance, the ESRP data shows that multiple disk subsystems already exceed this performance requirement. In addition, SPC-1 benchmark result trends indicate that the application performance gap is widening.

To date these higher performance requirements have most likely been met by over-provisioning of disk drives. However, on a price equivalent basis, at least in the case of TMS, flash SSS storage can provide a more economical solution.

Summary

Storage benchmark history clearly shows that disk performance is steadily declining on a capacity basis, while email application performance requirements are also declining but not as fast. As such, an application performance gap has been created. Traditionally, over-provisioning was used to address this disk-application performance gap. However, given the costs of disk storage, current flash SSS storage may already be more economical. As disk device performance continues to decline, the price performance breakpoint will also decline and flash SSS should become the preferred storage technology for database oriented applications. Even further, flash SSS will gradually address more and more of mainstream application requirements and thereby gain significant market share.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.

Silverton Consulting
Strategy, Storage & Systems
P.O. Box 492|Broomfield, CO 80038-0492|
720-221-7270| http://www.silvertonconsulting.com