OBJECTIVE ANALYSIS
Semiconductor Market Research

OBJECTIVE ANALYSIS COMMENTARY

# MAKING THE MOST OF FLASH STORAGE

Today's data center managers are faced with some difficult challenges: While data needs are ballooning at a 50% annual rate, the budget for IT is growing significantly more slowly at roughly 3% per year.

Meanwhile processor and DRAM speeds are increasing thanks to Moore's Law advances, while the speed of conventional storage (HDD, optical, and tape) is barely accelerating even though the storage per platter on HDDs is increasing at a rate to match data capacity needs.

On top of this data centers are beginning to widely embrace virtualization, and system architectures are changing to meet the needs of a virtualized data center with storage centralized in a SAN while the servers are gathered into clusters of identical hardware.

This creates an important problem: A growing gap has developed between the speeds of servers and the speed of the SAN providing data to those servers.

Fortunately all these changes are coinciding with the widespread availability of inexpensive NAND flash memory. Flash, often in the form of SSDs, has been found to match the needs of today's systems by providing a new layer in the storage hierarchy that is faster than HDD yet slower than DRAM, at a cost per gigabyte that is lower than DRAM, but higher than HDD.

Many systems already use flash to run their CRM, ERM, DBR, and Exchange systems faster.

The IT community no longer asks whether or not flash belongs in the data center – it clearly belongs. The questions today are:

1) How much flash to use?
2) Where to use it?
3) What data should be put into it?

Let's examine each of these questions.

## *How Much Flash to Use?*

The amount of flash a system needs depends to a great degree upon how it is managed. As we will see in a subsequent section considerably less flash is required if it is managed properly. This is because most stored data does not need fast access, only the data that is being used at the moment (Figure 1.) If only this data is stored in flash, then the vast majority of the data can be stored in less expensive HDDs.

**Figure 1. Not All Data is Regularly Accessed**

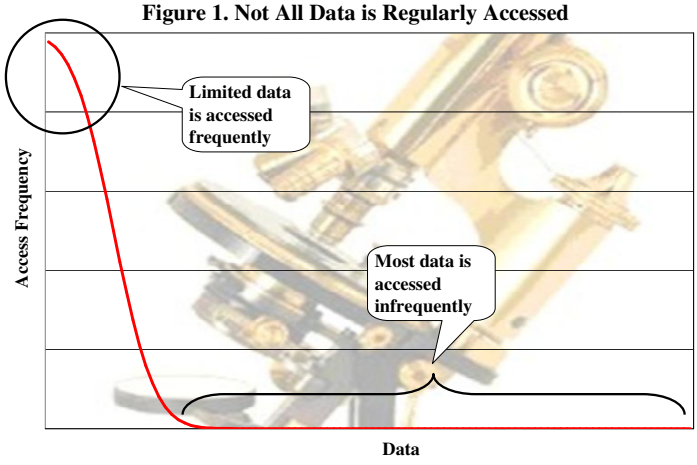


Although the minimum amount of flash needed is usually determined by trial and error, most users find that a nominal

flash investment boosts performance well beyond their performance goals. A rule of thumb is to use about 10% as many gigabytes of flash as there are of disk space, and to manage the flash with caching software. In many systems this will provide more than enough performance, but in some it may be less than the optimal amount. Providers are beginning to introduce tools that help estimate the flash needs of a particular environment.

## Where to Use Flash?

Flash can accelerate a system's performance in different ways, each of which is best suited to certain kinds of systems.

In systems where key data is not shared, and can be safely lost or discarded, the highest performance approach is to insert a large capacity flash card within the server on its PCIe channel, bringing the flash as close to the processor as possible. This approach minimizes latency to eliminate both hardware and software bottlenecks between the storage and the CPU.

**Figure 2. Performance Depends on Data Placement**

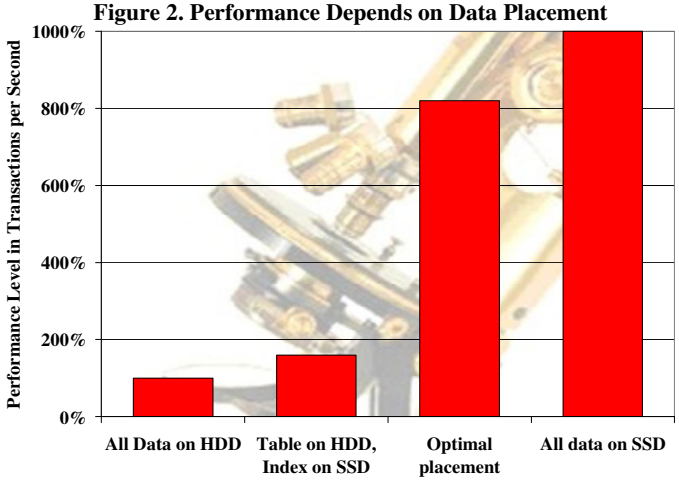


High availability (HA) systems that cannot tolerate data loss often use either a RAID or a SAN to protect the data. Shared data systems also use such configurations.

A good way to accelerate SAN storage is to cache each server's network I/O requests on fast flash storage within the server. An SSD or PCIe flash card is managed by caching software that replicates data in the cache as it is read from the SAN.

The performance of any RAID can be substantially improved through the use of flash, and if the RAID incorporates a modest amount of flash managed by caching software, its performance will approach that of an all-flash RAID. Typically this arrangement incorporates one or more high performance SSDs, but recently one vendor, LSI, has introduced a RAID card with internal flash to bring this same architecture to budget-oriented systems in small and medium-sized businesses.

## What Goes into the Flash?

Although the fastest solution would involve replacing the entire storage array with flash, this is rarely financially feasible. If we assume that the frequently accessed data can be stored in flash while the rest is stored on HDDs this brings us to a sticky situation – how do you manage this for the best performance?

There are two ways of managing a small flash backed by larger HDD storage. These are referred to as Manual Data Placement, and Automatic Data Placement.

Manual Data Placement requires an operator to decide which data is most likely to be frequently accessed and to move this data to the faster storage. This typically might include indexes, roots, journal files, some programs, and smaller databases. Such an approach is labor intensive as it requires hand tuning on an ongoing basis.

Automatic Data Placement allows the computer to measure which files (or

even blocks) are most frequently accessed, moving or copying those elements to the faster flash storage as required.

Although it is not intuitive, there is a significant advantage to using automatic data placement. Figure 2 illustrates this: researchers at IBM configured a system first using HDDs for all of the storage (left column) then SSDs (right column.) In this particular system the all-SSD system performed ten times as fast as the all-HDD system. They then moved the database index to SSDs leaving the database table on HDDs to achieve a 2x performance improvement over the all-HDD system. When this team implemented a caching algorithm using the same small pool of SSDs the performance of the mixed HDD/SSD system jumped to 8x the performance of the all-HDD system, or 80% of the performance of the all-SSD system at a fraction of the cost.

This makes it clear that caching algorithms not only reduce operator workload but also provide a significant performance boost over systems that use Manual Data Placement. The quality of the caching algorithm can provide further improvements, bringing the cached system's performance even closer to that of the all-SSD system. The right caching system guarantees flash-like performance using the right amount of flash.

## *Choosing a Partner*

Today there are estimated to be over 200 companies producing SSDs, and over a dozen manufacturers of PCIe NAND storage. Flash cache management software is also provided by over a dozen firms. How can an IT professional choose one over the other?

Although many believe that making a PCIe SSD involves little more than soldering together some NAND flash chips, a few off-the-shelf SATA SSD controllers, and a PCIe RAID controller, the optimum solution requires a significant engineering effort. The system must be designed for both high performance and ultimate reliability. The ideal provider would have a track record in this space, and would be known for not only the quality of its product but for the quality and reliability of its support.

This is of particular concern to the IT manager whose job is to assure that the data center continues to operate non-stop even when undergoing upgrades.

Of the companies that produce SSD controllers and RAID controllers, some have names that stand out more than others. The SSD controller market has been in flux – Jmicron held the leading position for a short while until displaced by Indilinx. Indilinx was in turn displaced by SandForce (acquired in early 2012 by LSI) which has successfully maintained performance leadership for two generations – quite a feat!

The RAID controller card market is a little less dynamic. LSI and Adaptec continue to hold the leading positions in this market as they have for a number of years.

Objective Analysis sees LSI's acquisition of SandForce as a very positive move for both companies and for their customers, since the companies have highly complementary products with a reputation for excellence. There are synergies in both technology and business that should make the fit turn into a very good endeavor.

***Jim Handy, June 2012***