

The Value of Flash for SMB Application Acceleration

Report prepared under contract with Hewlett Packard Enterprise (HPE)

The Benefits of Flash for Data Storage

Flash storage is computer storage that uses non-volatile memory technology (NAND flash) for the storage media rather than traditional magnetic media such as hard disk drives (HDD) or tape drives. Flash storage has revolutionized the IT industry and is now deployed in everything from small mobile devices to data center equipment found in large enterprises. Large enterprises have strongly embraced the use of flash-based media in data center storage over the last year or two because of its tremendous performance benefits, lack of moving parts, low noise, low weight, smaller footprint and reduced power and cooling needs.

In the past, flash-based media was prohibitive for most businesses due to cost. However, as pricing for flash technology has continued to drop, it has reached the point where deploying flash capacity is surprisingly affordable—especially in small amounts. This is the concept behind entry hybrid flash storage systems where flash is deployed alongside HDD technology.

Additionally, with flash storage, there are new metrics that need to be considered. Price per gigabyte for raw capacity should not be the only adoption criterion. Price per IOPS and IOPS per watt also need to be examined. For some time, we at Demartek have been saying that **“flash storage is dollars per gigabyte and pennies per IOPS while hard disk drives are pennies per gigabyte and dollars per IOPS.”** With the prices of flash dropping, the cost of flash is no longer dollars per

gigabyte. When considering IOPS per watt, flash storage is at least 100x better than traditional hard disk drive technology.

Flash technology in a hybrid storage system can be used to accelerate application performance in several ways, including as a *read cache* or as a *performance tier*. In addition to the obvious performance benefits, flash storage consumes lower amounts of power and provides more consistent performance than hard disk drive technology. Flash storage can also handle the demands of multiple workloads simultaneously.

One of the key benefits of flash storage is that it provides lower (better) latency. Increasingly, application latency is an important storage performance metric, often more so than IOPS or throughput. Lower latency allows transactional workloads to deliver much more consistent response to users, allowing them to continue their workflows without disruption.

When users see a “snappier” response from their applications, they sometimes believe that their computer systems have been upgraded. Upgrading the storage with flash-based media frequently obtains this benefit.

Once the decision is made to purchase a low-cost hybrid array or add flash technology to an existing storage system, the next choice is to decide how to use that capacity. Each of the two data placement models—read caching and performance tiering—has its own advantages and each will accelerate performance differently.

Flash or SSD Caching

When flash is used as a caching layer in front of hard disk drives, the following observations can be made:

- > Multiple applications can automatically gain a performance benefit.
- > Applications do not need to be modified to take advantage of the SSD cache.
- > The caching solution places a copy of “hot” data into the cache.
- > The caching solution decides when to place the copy of the hot data into the cache.
- > The aggregate performance gains occur over time as the cache “warms-up.”
- > Management is relatively simple.

Performance Tiering

Performance tiering takes the caching technology one level further by moving frequently accessed or “hot” data onto SSDs as the workloads access their data. This requires intelligent software and is considered a more sophisticated solution. Overall application performance is therefore superior for many if not most workloads, as performance tiers apply to both reads and writes. Key considerations to keep in mind are:

- > A performance tiering solution re-evaluates the I/O activity of the storage system on an ongoing or periodic basis and dynamically moves data to the SSD layer.
- > Data that is accessed frequently remains on the SSD tier, and data accessed less frequently is moved to the slower hard disk drive tier.
- > A performance tiering solution has many of the same benefits as read caching, with the added bonus of accelerating writes.
- > A good performance tiering solution will automate as much of this process as possible.

Typical Results

At Demartek, we have observed noticeable performance gains by deploying flash storage on a variety of workloads, including databases, server virtualization, operating system paging files and more.

Read caching is a solution that is particularly beneficial to read-heavy workloads that have data “hot spots” resulting from repeated accesses of the same data over various periods of time. Database applications frequently see performance benefits with read caching.

Performance tiering is a good fit for those applications that perform mixed read and write activity, which can include email servers, file servers, operating system paging files and collaboration environments. Database performance is often improved by placing indexes and temporary spaces onto a performance tier that accelerates reads and writes.

In one recent example, we ran tests using a hard disk drive storage system. Then we upgraded that storage system with flash read caching. Finally, we applied performance tiering software to create a flash tier. We used these three configurations to run the same performance and read-intensive database workloads based on Microsoft SQL Server.

The **read caching results** were as follows:

- > Database Transactions per second: +35%
- > Latency: 4.5x improvement
- > Throughput: 2.5x improvement
- > IOPS: 3x improvement

The **performance tiering results** were as follows:

- > Database Transactions per second: +60%
- > Latency: 11x improvement
- > Throughput: 4.6x improvement
- > IOPS: 4.5x improvement

Similar improvements can be seen in point-of-sale applications, medical office applications, knowledge worker environments and more.

The Demartek HPE MSA 2042 Performance Report:
www.demartek.com/HPEMSA2042

General Guidance

It doesn't take much flash storage to provide significant performance benefits to a hard disk drive storage system. We typically see very good improvements with flash storage comprising 3% - 5% of the total storage capacity of the storage system.

For those who want to get started slowly with flash storage, a good way to begin is with a hard disk drive array that can be upgraded to either a flash caching or flash performance tiering solution non-disruptively. It is important to choose a storage solution that provides affordable options to add flash in small increments that accelerates application workloads.

Regardless of size, even if a business or organization does not choose to include flash in their storage solutions today, they would be wise to use the ability to add flash at a future date as a key purchasing criterion.

We often find that once people embrace flash storage and see the benefits, they move from pilot projects to wholesale adoption of flash storage for as many workloads as possible.

Conclusion

Storage solutions that include some flash technology have become quite affordable recently, giving small and medium-size businesses new options for deploying flash to satisfy performance-hungry applications.

HPE has a full line of flash-enabled storage systems from the entry-level to the high end that complement the HPE ProLiant line of servers.

HPE has recently introduced two products, the HPE MSA 2042 and the HPE StoreVirtual 3200 that combine the benefits of flash technology with a low-cost, next-generation architecture that can accelerate application performance. The MSA 2042 comes with 800 GB of flash built-in for immediate application acceleration. The StoreVirtual 3200 offers flash as an option that can be added when the customer is ready.

These solutions can start as hard drive-only with the option to add flash storage non-disruptively in the future or can be configured with some flash storage from the factory. In addition, the HPE MSA 2042 includes the option to use built-in flash for either read caching or performance tiering without the need to purchase additional software licenses. This provides ultimate flexibility for businesses that know they want to accelerate application workloads from day one.

The most current version of this report is available at http://www.demartek.com/Demartek_Commentary_Value_of_Flash_for_SMB_Application_Acceleration_2016-08.html on the Demartek website.

Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.