



NETAPP WHITE PAPER

# Flash Memory Technology in Enterprise Storage Flexible Choices to Optimize Performance

Mark Woods and Amit Shah, NetApp  
November 2008 | WP-7061-1008

## EXECUTIVE SUMMARY

Solid state drives (SSDs) based on flash memory are generating a lot of excitement. This enthusiasm is warranted because flash SSDs demonstrate latencies that are at least 10 times lower than the fastest hard disk drives (HDDs), often enabling response times more than 10X faster. For random read workloads, SSDs may deliver the I/O throughput of 30 or more HDDs while consuming significantly less power per disk. The performance of SSDs can reduce the number of fast-spinning hard disk drives you need in a storage system. Fewer disk drives translates into significant savings of power, cooling, and data center space. This performance benefit comes at a premium; flash SSDs are far more expensive per gigabyte of capacity than HDDs. Therefore SSDs are best applied in situations that require the highest performance.

The underlying flash memory technology used by SSDs has many advantages, particularly in comparison to DRAM. In addition to storage persistence, these advantages include higher density, lower power consumption, and lower cost per gigabyte. Because of these unique characteristics, NetApp is focusing on the targeted use of flash memory in storage systems and within your storage infrastructure in ways that can deliver the most performance acceleration for the minimum investment.

We are implementing flash memory solutions using SSDs for persistent storage, and we will also use flash memory directly to create expanded read caching devices. Caching can deliver performance that is comparable to or better than SSDs. Because you can complement a large amount of hard disk capacity with a relatively modest amount of read cache, caching is more cost effective for typical enterprise applications. As a result, more people can benefit from the performance acceleration achievable with flash technology.

You get even more flexibility and value from flash technology by combining it with the NetApp® unified storage architecture, which enables you to leverage your investment in flash memory to simultaneously accelerate multiple applications, whether they use SAN or NAS. Storage efficiency features such as deduplication for primary storage further increase your power, cooling, and space savings.

This white paper is an overview of NetApp's plan to deliver SSDs (both native and virtualized arrays) plus flash-based read caching and of our ability to further leverage both of these technologies in caching architectures. Selection guidelines are provided to help you choose the right technology to reduce latency and increase your transaction rate while taking into consideration cost versus benefit.

## TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION .....</b>	<b>3</b>
	<b>USING FLASH TECHNOLOGY FOR ENTERPRISE STORAGE .....</b>	<b>3</b>
<b>2</b>	<b>PERSISTENT STORAGE WITH SOLID STATE DRIVES .....</b>	<b>4</b>
	<b>WHERE SOLID STATE DISKS FIT IN ENTERPRISE STORAGE .....</b>	<b>5</b>
	<b>NETAPP SSD PLANS .....</b>	<b>5</b>
<b>3</b>	<b>READ CACHING .....</b>	<b>5</b>
	<b>NETAPP READ CACHING PLANS .....</b>	<b>5</b>
<b>4</b>	<b>CACHING ARCHITECTURE .....</b>	<b>7</b>
	<b>NETAPP CACHING ARCHITECTURE OPTIONS .....</b>	<b>8</b>
<b>5</b>	<b>CHOOSING THE RIGHT TECHNOLOGY FOR YOUR APPLICATION .....</b>	<b>8</b>
<b>6</b>	<b>MULTIPLYING THE BENEFITS OF FLASH .....</b>	<b>9</b>
<b>7</b>	<b>CONCLUSION: WHAT YOU CAN EXPECT FROM NETAPP .....</b>	<b>10</b>

## 1 INTRODUCTION

Solid-state disks (SSDs) based on flash-memory technology have gotten a lot of attention recently. After years of development for use in the consumer market, SSDs are reaching both a maturity and a price point where enterprises can consider them seriously.

The standard hard disk drives (HDDs) that have been the mainstay of enterprise storage for decades have some nagging deficiencies that make provisioning storage for applications with high I/O demand difficult. Today's fastest electromechanical disk drives have access times of 3-4 milliseconds\*, which sets a lower bound on read response times.

Further, HDDs are capable of performing no more than 300-400 random I/O operations per second (IOPS). As a result, a storage system capable of handling tens of thousands of IOPS requires hundreds of spindles to achieve the necessary level of performance—even in situations where the capacity is not needed. Hundreds of disks in turn consume significant rack space, power, and cooling. Flash-based SSDs, by comparison, achieve latencies 10-20 times lower and can perform tens of thousands of IOPS per device.

The conventional wisdom is that because of their performance advantages, coupled with significantly lower power consumption, SSDs will supplant the high-performance HDD market over time. Industry analysts believe that SSDs will be deployed as an ultra-high performance storage tier within networked storage (which IDC defines as tier 0) and slowly cannibalize the market for high-performance HDDs as the cost continues to fall.

Although NetApp agrees with this assessment, we are taking a broader approach to flash technology. We recognize that it is the unique characteristics of flash memory—including high density, low power consumption, and relative economy in comparison to DRAM—that are driving the adoption of flash-based SSDs.

Our strategy is to enable you to leverage these flash advantages more quickly and for a wider variety of applications. We will combine flash technology with our unified storage architecture in innovative ways to meet the performance needs of a range of applications.

We'll achieve this by providing both flash-based SSDs for persistent storage and flash-based caching devices that will deliver the performance benefits of flash memory without the cost and complexity of an additional storage tier.

### USING FLASH TECHNOLOGY FOR ENTERPRISE STORAGE

Although flash technology offers significant performance advantages over even the fastest HDDs, that performance comes at a significant price premium per gigabyte of storage. HDDs typically cost \$1-\$2 per gigabyte, while the cost of comparable SSDs currently ranges from \$15-\$100 per gigabyte. As a result, you must use flash for the workloads where it gives you a high return on investment for the performance gain to justify the added expense.

Flash technology—whether deployed in an SSD or as a cache—provides maximum advantage for applications that perform many random read I/Os. These are the applications that get the greatest benefit from the improved latency of flash memory versus HDDs. With a minimum 10X improvement in latency versus HDDs, flash technology yields a noticeable performance improvement for interactive applications and a reduced execution time for compute-intensive applications.

Applications that primarily generate random read I/O can push HDDs to their limits, forcing disk heads to seek before each operation. When you see published statements suggesting that a single flash SSD can deliver the performance of 30 HDDs, those numbers refer to random read workloads such as those generated by OLTP and similar applications. (We'll discuss specific applications in section 5.)

---

\* All HDD performance numbers in this paper are based on the published performance of the Seagate Cheetah 15K.6 drive, a recently released 15K RPM, Fibre Channel disk drive. Similarly, SSD performance numbers are from the published performance numbers for the STEC Zeus<sup>IOPS</sup> flash-based solid-state drive. Both are considered to be among the best in their class.

For applications with sequential workloads—where large chunks of data are read at one time—the relative advantage of flash technology is diminished. This is because standard HDDs are much better at sequential workloads since disk heads have to be repositioned much less frequently; HDD and SSD transfer rates are comparable.

Similarly, the advantage of flash technology over HDDs is not as great for writing. Because of the way flash memory writes data, it delivers lower write performance than read performance. In many circumstances, HDDs deliver extremely good write performance because they are able to write large chunks of data without repositioning the disk heads. This is particularly true with modern file systems such as WAFL® that buffer write data and perform sequential writes..

This paper explores the range of options that NetApp will deliver to help you leverage the unique performance advantages of flash memory for applications that perform intense random read I/O. These options include:

- **Persistent storage using SSDs**
- **Read caching**
- **A caching architecture capable of using both SSDs and read caching**

Guidelines are provided for choosing between these options for different applications, and the advantages of deploying these technologies in combination with the NetApp unified storage architecture are discussed.

## 2 PERSISTENT STORAGE WITH SOLID STATE DRIVES

The performance of SSDs is compelling for certain I/O-bound applications. A flash-based SSD achieves access times measured in microseconds and can perform tens of thousands of IOPS.

Table 1) Comparison of a 15K RPM hard drive to a flash-based SSD.

	Seagate Cheetah 15K.6 HDD	STEC Zeus <sup>IOPS</sup> SSD
<b>Access Time</b>	3-4 milliseconds <sup>1</sup>	20-120 microseconds
<b>Random Read Performance</b>	300-400 IOPS <sup>2</sup>	45,000 IOPS
<b>Capacity</b>	300GB, 456GB	256GB, 512GB <sup>3</sup>
<b>MTBF</b>	1.6 million hours	NA <sup>4</sup>
<b>Power (Watts)</b>		
<b>Idle Mode</b>	12	5.4
<b>Operating Mode</b>	17	8.4
<b>Estimated Cost<sup>5</sup></b>		
<b>\$/GB</b>	\$1-\$2/GB	\$65/GB
<b>\$/IOP</b>	\$1-\$2/IOP	\$0.20/IOP

<sup>1</sup> Seek time

<sup>4</sup> Manufacturer claims 2x HDD MTBF

<sup>2</sup> Typical for 15K RPM HDD

<sup>5</sup> Based on available Internet pricing

<sup>3</sup> Unformatted

Although performance is a clear advantage, there are obviously other factors to consider with SSDs, such as capacity and reliability. Capacity is increasing rapidly; although capacities don't equal those available for SATA disk drives, they are comparable to the available capacities for the fastest Fibre Channel drives.

Reliability of flash-based SSDs has been a concern because each block can be erased and rewritten only so many times (currently 100,000 write cycles for SLC NAND flash). However, wear-leveling algorithms mean that a drive's capacity is used evenly to avoid failures and increase longevity. Manufacturers are combining wear leveling with spare capacity and caching technologies for further SSD endurance under random write workloads. Mean Time Between Failures (MTBF) for enterprise-class SSDs is now said to exceed that for electromechanical disks, and SSDs may be less susceptible to shock, vibration, and catastrophic failure.

Given that capacity and reliability concerns have been addressed, cost remains the biggest barrier to the adoption of SSDs.

## WHERE SOLID STATE DISKS FIT IN ENTERPRISE STORAGE

For now, SSDs are best targeted to applications where performance is the most important consideration. Early adopters are creating a new storage tier composed of SSDs and targeted for hot data. To effectively leverage such a storage tier, automated migration software may be needed to move data between tiers as usage changes. Barring that, a lot of administrator time may be required to manually optimize the use of this high-performance tier. A caching architecture may offer an alternative to eliminate the data movement needed with a tier 0 SSD deployment. (See section 4.)

## NETAPP SSD PLANS

For SSD storage, NetApp will provide two options:

- SSDs in standard NetApp disk shelves
- The ability to virtualize third-party SSD disk arrays

NetApp is working to qualify the best available enterprise-class SSDs in a range of capacities. You'll be able to deploy this storage tier in your existing NetApp storage systems alongside existing Fibre Channel and SATA disk storage.

In addition, NetApp will provide the ability to virtualize third-party SSD arrays by using NetApp V-Series storage systems, allowing you to combine the performance benefits of optimized SSD storage arrays with the full range of NetApp unified storage capabilities.

## 3 READ CACHING

Given the high cost per capacity of current SSDs, read caching is an alternative that may be better suited to a broader range of enterprise applications. Like SSDs, read caching can reduce response times and accelerate throughput, while cutting the total number of HDDs needed to achieve a given performance level. Given the relative economy of flash memory in comparison to DRAM, a large secondary flash read cache could supplement a storage system's existing cache memory to substantially increase cache hits and avoid disk reads altogether.

Like SSDs, read caching offers the most benefit for applications with a lot of small, random read I/Os. Once a cache is populated, it can substantially decrease the average response time for read operations and reduce the total number of HDDs needed to meet a given I/O requirement. Reducing disk count in turn translates into savings on power, cooling, and data center space.

Because a read cache resides in the data path rather than in the storage stack, it can deliver even faster response times than SSDs. You also gain substantial performance improvement with far fewer gigabytes of total flash memory, so adding cache can be substantially cheaper than adding a separate SSD storage tier. This makes read caching feasible for a broader range of applications. Because caching is dynamic, you don't need any special software or manual data management to make sure that the cache contains your "hot" data.

## NETAPP READ CACHING PLANS

NetApp is pursuing a strategy of offering modular read caches that can be installed in many existing NetApp systems. The Performance Acceleration Module (PAM), a DRAM-based intelligent read cache that began shipping in mid-2008, provides 16GB of extended cache memory per module. Up to five modules are supported for a maximum of 80GB of cache depending on storage system model and available PCIe expansion slots. The caching policies implemented for PAM are designed to distinguish high-value, randomly read data from sequentially read data and maintain that data in cache to avoid time-consuming disk reads.

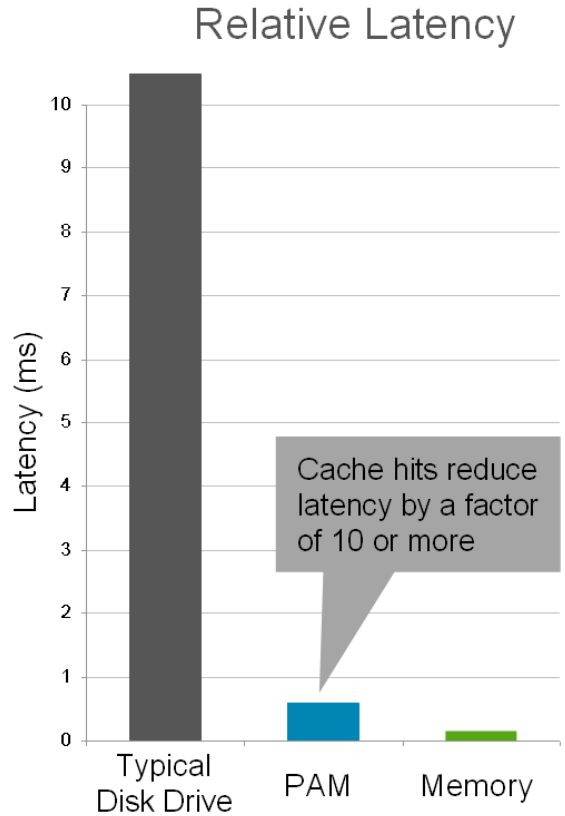



Figure 1) NetApp Performance Acceleration Module (PAM) reduces latency for random reads.

A second-generation, flash-based Performance Acceleration Module is under development. This module will benefit from the greater density of flash memory to offer more cache memory per module. As a result, the working set of large applications such as databases can be accommodated entirely in cache.

NetApp has also developed Predictive Cache Statistics (PCS) software that can be run on your existing storage systems to determine whether your workloads will benefit from read caching and, if so, how much cache you will need.

## 4 CACHING ARCHITECTURE

A third flash-based alternative—creating a caching architecture within your storage infrastructure—builds on the previous two approaches. With a caching architecture, a small, fast storage cache resides logically between your primary storage systems and your compute servers or desktop clients. As with an internal read cache, data is automatically copied into your caching tier the first time it is read; subsequent reads are satisfied from the cache rather than from the origin storage system. By concentrating your investments in flash technology in a centralized caching tier, you can leverage that investment across multiple storage systems and applications. You can also deploy more economical, high-capacity disk drives in your primary storage without adversely affecting performance.

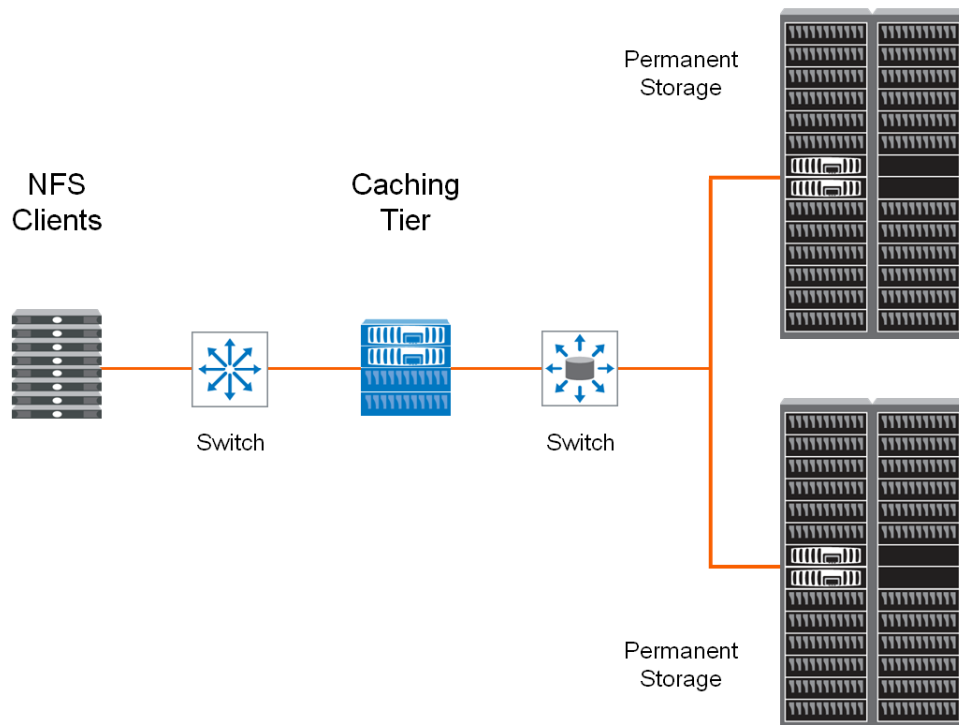


Figure 2) By concentrating flash investments (cache and/or SSDs) in a caching tier, the resource can be leveraged by multiple storage systems. Hot data is automatically cached, so manual data management or migration software is not needed.

This approach has already demonstrated significant benefits for accelerating the speed of parallel software builds as well as compute-intensive applications such as movie rendering, electronic design automation, seismic analysis, and financial market simulation. Any application in which the same data is read in parallel by multiple compute servers can benefit from such an approach.

## NETAPP CACHING ARCHITECTURE OPTIONS

NetApp has pioneered caching architectures to accelerate NFS storage performance with its FlexCache™ software and storage acceleration appliances. FlexCache eliminates storage bottlenecks without requiring additional administrative overhead for data placement. The addition of SSDs and/or read cache to a NetApp caching architecture will extend the benefits of FlexCache by reducing response times for critical applications, boosting transaction rates, and cutting energy consumption in your caching layer. Like any storage system, the I/O performance of a caching appliance depends on the number of disk spindles. Flash-memory technologies—either SSDs or cache—will significantly reduce the number of disks needed to achieve a given performance level while also reducing response time.

NetApp will make its full range of flash technology options—native SSDs, virtualized SSD arrays, and read cache—available for use in its caching architecture. Active data is automatically cached, giving you the performance benefit of tier 0 storage for repeated reads without requiring automated data migration software or manual data management. If you optimize your caching tier for performance, you can optimize your permanent storage tier for cost at the same time without adverse impacts.

## 5 CHOOSING THE RIGHT TECHNOLOGY FOR YOUR APPLICATION

As noted earlier, any application that generates a significant amount of random read activity may experience a performance benefit from the use of flash technology. The industry is still in a learning mode, so the guidelines for each application aren't chiseled in stone, but here are some rules of thumb.

General considerations:

- Consider the relative total cost of the available alternatives. This should include upfront costs and ongoing administration. Will the alternative deliver a reasonable return on investment (ROI) in terms of the expected performance gain, space savings, and power and cooling savings?
- What risks are associated with the deployment? How disruptive will it be to ongoing operations?

Specific guidelines:

- Flash SSDs are currently expensive per gigabyte of capacity and are best suited for applications where reducing latency and/or achieving high IOPs is paramount. For example, financial modeling and other HPC applications such as electronic design automation, computer-aided engineering, and seismic data analysis typically yield a high ROI for performance improvements. In some instances, critical enterprise applications—overloaded OLTP applications, Web serving, and so on—may also benefit from targeted use of SSDs.
- Use flash memory as a read cache in the storage controller as a cost-effective way to optimize performance for common enterprise applications such as file services, Microsoft® Exchange or other e-mail, Web hosting, and OLTP databases. NetApp Predictive Cache Statistics is available to help you determine whether your application can benefit from caching and to help you decide how much cache to purchase.
- Combine flash read cache and/or SSDs with FlexCache software to create a caching architecture for NFS environments. This caching tier may enable you to leverage your investment in expensive flash storage across multiple applications and storage systems. Because hot data is cached automatically, your caching tier always contains your working data set with no manual effort or additional software.

Of course there are exceptions to these guidelines based on your unique IO requirements.

Table 2) Pros and cons of available flash storage technologies.

Technology	Pros	Cons
<b>Solid State Drives</b>	<ul style="list-style-type: none"> <li>• Response times are consistently fast for reads</li> <li>• Low cost per IOP</li> <li>• Administrator has direct control over data stored in SSD tier</li> </ul>	<ul style="list-style-type: none"> <li>• High cost per gigabyte</li> <li>• Requires software tools and administration to move hot data into and out of SSD tier</li> </ul>
<b>Read Cache</b>	<ul style="list-style-type: none"> <li>• Hot data automatically flows into read cache—no administration required</li> <li>• Deployment is relatively nondisruptive</li> <li>• Viable for common enterprise applications</li> </ul>	<ul style="list-style-type: none"> <li>• Cache must be populated before it becomes effective</li> </ul>
<b>Caching Architecture with SSDs or Read Cache</b>	<ul style="list-style-type: none"> <li>• Hot data automatically flows into the caching tier</li> <li>• Deployment is relatively nondisruptive</li> <li>• Scalable solution for high-performance applications</li> </ul>	<ul style="list-style-type: none"> <li>• Cache must be populated before it becomes effective</li> <li>• NetApp FlexCache software currently supports only the NFS storage protocol</li> </ul>

## 6 MULTIPLYING THE BENEFITS OF FLASH

Like any IT investment, you need to get the most from your investment in flash technology, no matter which approach you choose. NetApp storage systems combined with the NetApp Data ONTAP® storage operating environment are uniquely designed to enable you to use these technologies to their maximum extent. Many of your existing NetApp storage systems will be able to take advantage of flash technology. NetApp won't limit availability to just a few high-end systems.

Data ONTAP includes a number of unique features to help you get the most from all your storage—including flash memory. Whether you choose SSDs or a read cache, Data ONTAP provides features to significantly improve storage utilization. For instance, NetApp deduplication improves storage efficiency by eliminating duplicate blocks on primary storage to help you get the most out of your investment in both SSDs and read cache. (Deduplication of primary storage means that identical data blocks aren't cached in multiple locations in your read cache.) Using deduplication, thin provisioning, and other NetApp storage efficiency features in combination with flash memory can multiply your savings in terms of power, cooling, and space.

These features are a standard part of the NetApp unified storage architecture. With unified storage, a single storage system can meet all your NAS and SAN storage needs. Rather than provisioning SSDs for separate NAS and SAN deployments, you can provision a single system with SSDs and/or cache to meet both needs and re-provision capacity for use by any protocol as requirements change. You can also virtualize third-party, optimized SSD disk arrays to gain the advantages of NetApp data management and multiprotocol support.

You can find out more about NetApp unified storage in a recent white paper, *Unified Storage Architecture: Enabling Today's Dynamic Data Center*<sup>†</sup>.

<sup>†</sup> <http://media.netapp.com/documents/wp-7054.pdf>

## 7 CONCLUSION: WHAT YOU CAN EXPECT FROM NETAPP

The case for flash memory technologies in the enterprise—both SSDs and read caching—is a strong one. NetApp has committed to support both these technologies, and you will be able to use them with many existing NetApp storage systems as well as in caching architectures with NetApp FlexCache. We will also support virtualized third-party SSD arrays with our V-Series open storage controllers.

Our goal is to provide you with more flash-technology options so that you can afford to use flash memory today to accelerate a range of applications, from the most performance-sensitive HPC applications to common enterprise applications such as e-mail, file service, and Web hosting. The combination of these technologies with the NetApp unified storage architecture provides great synergy, enabling you to get the most from your investments in flash technology.

NetApp will make a series of flash technology product announcements in 2009.

© 2008 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexCache, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft is a registered trademark of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7061-1008.