WHITE

PAPER

NetApp's Solid State Hierarchy

With a Focus on 'PAM II'

By Mark Peters

September, 2009

Table of Contents

Introduction

Solid state storage is an undeniably 'hot' topic these days. Two years ago, it barely warranted a conversation in general IT circles and now, it is hard to find a vendor or commentator that is not discussing it. What changed? Both supply and demand aspects have shifted markedly. The supply side has been driven by improved technologies, more attractive pricing, and the continuing inability of hard disk drives to keep pace with an increasingly wide 'storage-to-server' performance imbalance. On the demand side are increasingly IO-voracious applications and expectations, which can be extremely hard to meet, either:

a) *at all*—often forcing 'workarounds' such as short-stroking, and/or;
b) *efficiently*—the current focus on better use of resources demands new approaches.

As a result, solid state storage is set to be a mainstream IT technology. The decision to use it or not is based on business and economic value. Unfortunately, the enormous hype surrounding solid state storage can obscure its practical value and even lead to misunderstanding. Furthermore, because many vendors are implementing the technology simply by saying 'oh sure, you can put an SSD in my rack,' there's a lack of depth that goes into much of the discussion; as the old adage says, "if the only tool you have is a hammer, then everything looks like a nail." NetApp offers a wide range of solid state tools and is therefore open to a broad discussion around the value of solid state in various implementations. This paper looks at solid state in general before reviewing specific NetApp offerings and their applicability. The essence of solid state does not have to be overly complicated—it's not just a nerdy jaunt for the few, but a pragmatic business tool for the many. Ironically, it as much about capacity as it is about performance.

Solid State – Where and Why?

Solid State Storage, SSD, and the Potential for Value

It is a fact: the number one advantage of solid state is the improved performance that it offers. In ESG's latest annual Enterprise Storage Survey, a number of specific (flash-based SSD) questions were asked on the topic. Figure 1 shows both the most important considerations for users and the ranking of individual considerations (since respondents could have multiple responses). That users see value coming from higher performance, better reliability, and lower power and cooling requirements is probably no surprise. But some other aspects of solid state are less expected or obvious. For instance, solid state storage is not new—the first enterprise drive was shipped in 1978 and non-volatile flash first appeared in the 1980's, although only recently has it attained the capacities, reliability, data accuracy, and cost structure that make widespread commercial IT adoption practical. Indeed, there are still multiple media types—in addition to DRAM, the most common current media is NAND flash, whether delivered as Single Level Cell (SLC) or Multi Level Cell (MLC). The difference is the number of bits that can be stored in a given amount of media; to date, SLC is generally preferred for enterprise use since fewer cells has equated to better longevity and read/write cycles. However, MLC is also being packaged to suit enterprise use and there are many extended media technologies at the R & D or early product development phases. The most important point is that <u>solid state storage is not just solid state disks.</u> The three main current implementations are:

a. **'Solid State Disk'** – flash memory packaged into a standard disk form-factor

b. **'Flash Array'** – a high density, standalone solid state storage device

c. **Server Component** – built into a server as either 'expanded cost-effective memory' or storage

NetApp is working in all these areas. It will offer 'regular' SSD's built into its storage racks; it partners with Texas Memory Systems to deliver a large-scale SSD Array; and it offers the PAM (Performance Acceleration Module), which sits in the controller—a specialized server, with added functionality—and offers an expanded level two cache (more on this later) as opposed to regular persistent storage.

FIGURE 1. CONSIDERATIONS FOR USE OR INTEREST IN SOLID STATE DISK

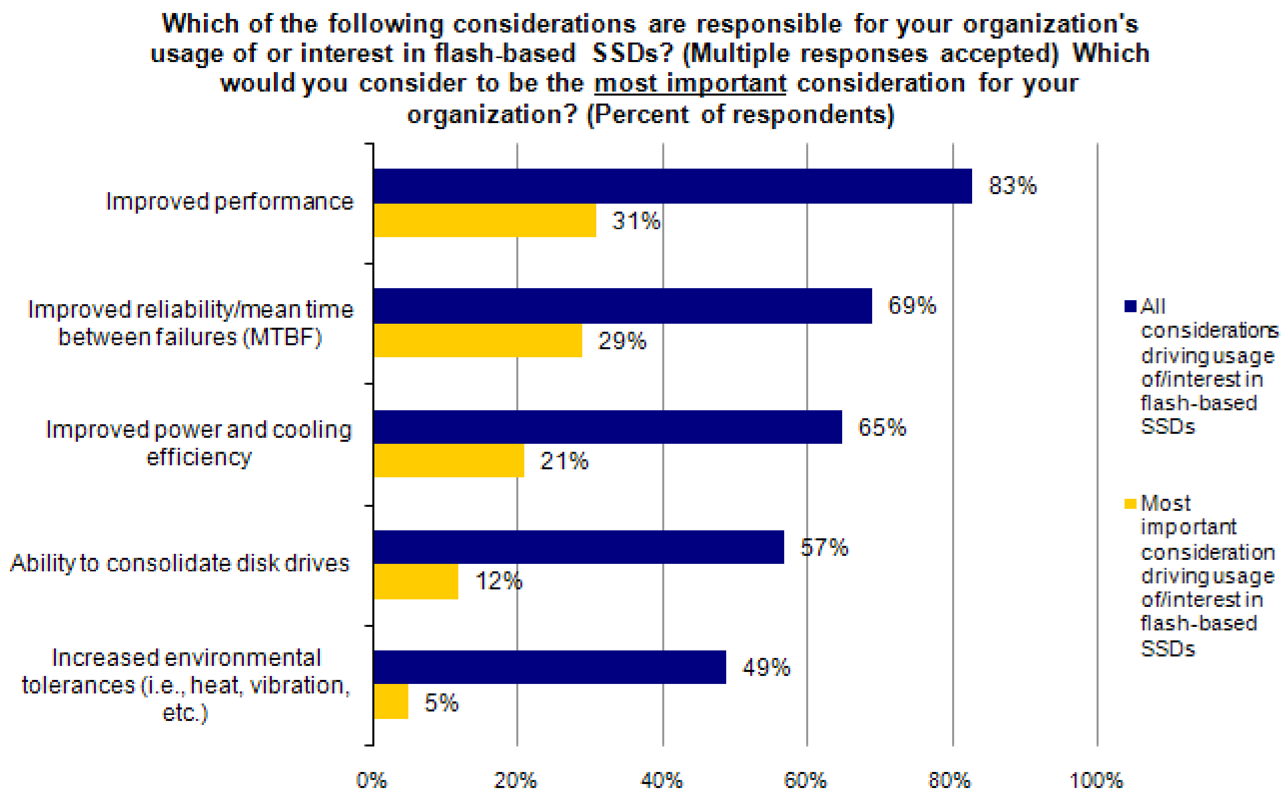

Source: Enterprise Strategy Group, 2009

Using all this performance—enhanced, as we've now seen, by flexibility of implementation—is both an opportunity and a challenge for IT managers. Certainly speed, fast response, and great throughput all sound good, but not at any price. Ultimately, the implementation must make business sense.

Opportunity: The *obvious opportunity* for IT management is clearly performance. Whereas RAM (main memory) response is measured in nanoseconds, that of HDDs is measured in milliseconds (3-4 at best, often 10-20). Solid state NAND flash sits in between the two and fills the performance gap with just microsecond response times. The *less obvious opportunity* is the effect that handling the most critical and intensive IO on solid state can have on other aspects of an operation—by reducing or precluding the need for 'unnatural' acts (such as short stroking) in order to achieve the requisite IOPS, it is actually possible to save money; both by needing less spindles and by needing fewer of the fastest disks drives.

Challenge: The *challenge* for IT management is to learn to think differently. Most users have grown up in a world where the storage value metric has been cost per GB. This still tends to be used, even when data centers do things like short-stroking, which means they are really buying (and thus limited to) the IO capacity of that particular drive. Flash memory requires a new set of metrics. If you are buying performance, then performance should be measured. In other words, the measurement should be IOPS/$. After all, improving IOs is the reason to have solid state storage and is something that should be viewed as part of the traditional storage hierarchy.

Why Does the Storage Hierarchy Exist?

The storage hierarchy is purely and simply a pragmatic response to the economics involved with storage. Everyone makes cost/performance trade-offs about what storage tier makes economic sense for any given user, state-of-the business, or application. With this understood, it is easier to appreciate solid state storage—not just as a 'new' tier, but potentially as a number of tiers in the storage hierarchy depending on what solid state technology is used and where it is placed. Solid state technologies invariably deliver the optimum combination of the highest performance with lowest cost and power use. The only questions are, how much of a user's data deserves the 'top of the hierarchy' treatment and, at the same time, how much can be economically justified? Soon, and for the foreseeable future, this logic will lead us to a situation where high performance needs (IO) will most likely be best addressed by solid state storage and high capacity requirements will be addressed by mechanical devices.

Specific IT and Business Challenges

Key Application Areas and Deployment Models for Solid State Storage

Put simply, there are two main application areas for solid state memory: A*ny* IO intensive application can typically benefit from flash as can *any* application that demands the fastest possible response time across as much of its capacity as possible (i.e., as can be afforded). Thus, some typical use models might be for high performance computing, web services, e-mail, financial and telecom applications, and database acceleration. But the list is not exclusive—wherever performance is crucial and the volumes can justify the price, solid state storage can provide benefits.

The optimum deployment model depends on what users are trying to achieve. Clearly, persistent data should go on a flash-based storage device (a 'tier') whereas an overall storage performance boost can best be achieved by using flash in a cache/controller/appliance model. Whatever the implementation, one of the key benefits is that there is likely to be an associated improvement in other parts of the hierarchy—in other words, once the major IO activity (which is almost always from a small percentage of a user's overall capacity) is being served by solid state storage, the remaining, less active capacity can safely reside on less IO capable higher-capacity drives. Put simply, FC and fast HDDs can be replaced by slower, higher capacity, and more economical HDDs.

Solid State's Impact on IT and Business Challenges

Solid state storage represents a proven method to address some very pressing IT and business challenges.

IT Challenges: Data center managers are being asked to provide consistent, if not improving, levels of access to growing amounts of data. Doing so with standard tools is problematic because:

- Increases in data volumes far exceed any budget increases.
- The laws of physics dictate that contention only increases as HDD capacities increase.
- HDD performance has improved many times less than that of processors and applications.
- User expectations and needs demand ever-higher performance and system reliability.

Business Challenges: While IT demands are increasing, so is the pressure on data center managers to meet them at a lower overall cost with improved energy use. Terms such as 'green' and 'efficiency' are challenging IT professionals to evaluate new approaches. The demand for effectiveness ('get the job done') that has driven IT for decades is now joined by a demand for efficiency ('get the job done at least as well, and preferably better, than before, but do it using as few resources as possible').

Solid state storage devices actually offer the compelling ability to deliver simultaneously against both IT and business challenges. By using an optimum amount of capacity to serve the highest possible percentage of IO from a device that needs no 'unnatural acts' (such as short-stroking or over purchasing of spindles) in order to do so and is highly resource-efficient, solid state storage represents a potential gain in IT performance allied to better business metrics (whether those are derived from the economy, mandated, or moral obligations). As such, its benefits can be enjoyed across a wide spectrum of applications. ESG research has found that enterprises concur and see broad applicability for the technology; asked for their current perception of flash-based SSDs in enterprise storage systems, 12% already see the technology as having widespread applicability today and a further 61% see it as niche today with widespread applicability down the road.

Addressing Challenges with NetApp Solid State

When looking at the offerings NetApp has in the solid state arena, it is obvious that the company 'gets it.' Certainly, it has specific options today, but far more important is the fact that NetApp understands that solid state can be delivered—*should* be delivered—in multiple places within the storage hierarchy and the data infrastructure. Eschewing a slavish devotion to any one particular 'correct' method or media leaves NetApp open to develop and improve its offerings as the technologies advance, thus ensuring its users have both the latest products and also a range of options. As a quick reminder, NetApp currently offers –

- **Persistent data storage** options via either SSD in their racks (in the future) or the Texas Memory Systems RamSan-500 Flash SSD Array (today) behind a NetApp V-Series open storage controller. This array can offer up to 2 TB of capacity and deliver some 100,000 IOPS at around a 1 millisecond response time. Compared to a configuration of 15k RPM disk drives having the same capacity, this would be 10 times faster and roughly 15 times the IOPS. Alternatively, if users wanted to produce the same IOPS from either this SSD Array or 'regular' disks, they would find disk still roughly 10 times slower and would have to over-purchase capacity just to get the IO capability.

- **Cache extension with PAM II**; the SLC NAND flash-based PAM II (Performance Acceleration Module II) joins the DRAM-based PAM card NetApp introduced just a year or so ago. Each PAM II card is up to 512 GB, with up to 4 cards in PCI Express slots in the controller (meaning no additional rack space is consumed). PAM family modules are intelligent, level 2 read caches for the storage controller. They cache metadata and user data with the aim of avoiding disk reads and the added latency and limited IO throughput that result from such reads.

Solid State in the Storage Controller

Some potential for confusion exists in comparisons between the PAM implementation and other PCIe based solid state cards, most probably because IT managers are focusing on the PCIe form factor as opposed to functionality and context of use. Whether solid state technology is used as persistent storage or as cache, the architectural context matters—PAM is used in SAN and NAS environments, while typically other PCIe cards are used in a DAS (i.e, server attached) context. Another key distinction is functionality. PAM is a dedicated, intelligent read cache that is tightly integrated with the NetApp Data ONTAP operating system. Generally speaking, the server-based flash cards are used for persistent storage and are treated as block storage devices. There are advantages and uses for each approach, with the operational differences and key user implications being:

Server based, DAS focused solid state cards are good for:

- Cost constraint – depending on the deployment, server based storage can be less costly than networked storage
- CPU proximity – other than main memory, there is no faster way to get data to the CPU

Storage controller based, network (SAN/NAS) focused solid state cards are good for:

- Sharability – it is a practical need or an absolute requirement for many applications that data be shared
- Manageability – networked storage can be managed with greater flexibility
- Availability – networked storage generally has better HA (High Availability) characteristics
- Efficiency – Tools such as thin provisioning, snapshots, cloning, and deduplication are more available for (or limited to) networked implementations and increase usable capacity while keeping raw capacity constant

- Ease of backup, recovery, and disaster recovery – tending to be better and more rugged in networked environments

PAM Operational and Efficiency Impacts

The impact and value of PAM can be enormous. In a SPEC SFS 2008 File Services benchmark on a FAS3140, NetApp demonstrated an ability to drive the same IO throughput with fewer drives, improved response, and the option to use high capacity SATA drives for primary storage. CAPEX reductions were around 27% (due to fewer drives—FC or SATA), while OPEX was driven down by using 44% less space and 47% (FC) or 54% (SATA) less power. Driving down OPEX matters a lot in the real world. Figure 2 makes the point starkly: when asked which considerations would be most important in justifying IT investments over the next 1-2 years, 63% of respondents put OPEX reduction at the top of the list.

FIGURE 2. MAJOR CONSIDERATIONS DRIVING IT INVESTMENTS OVER THE NEXT 12-24 MONTHS

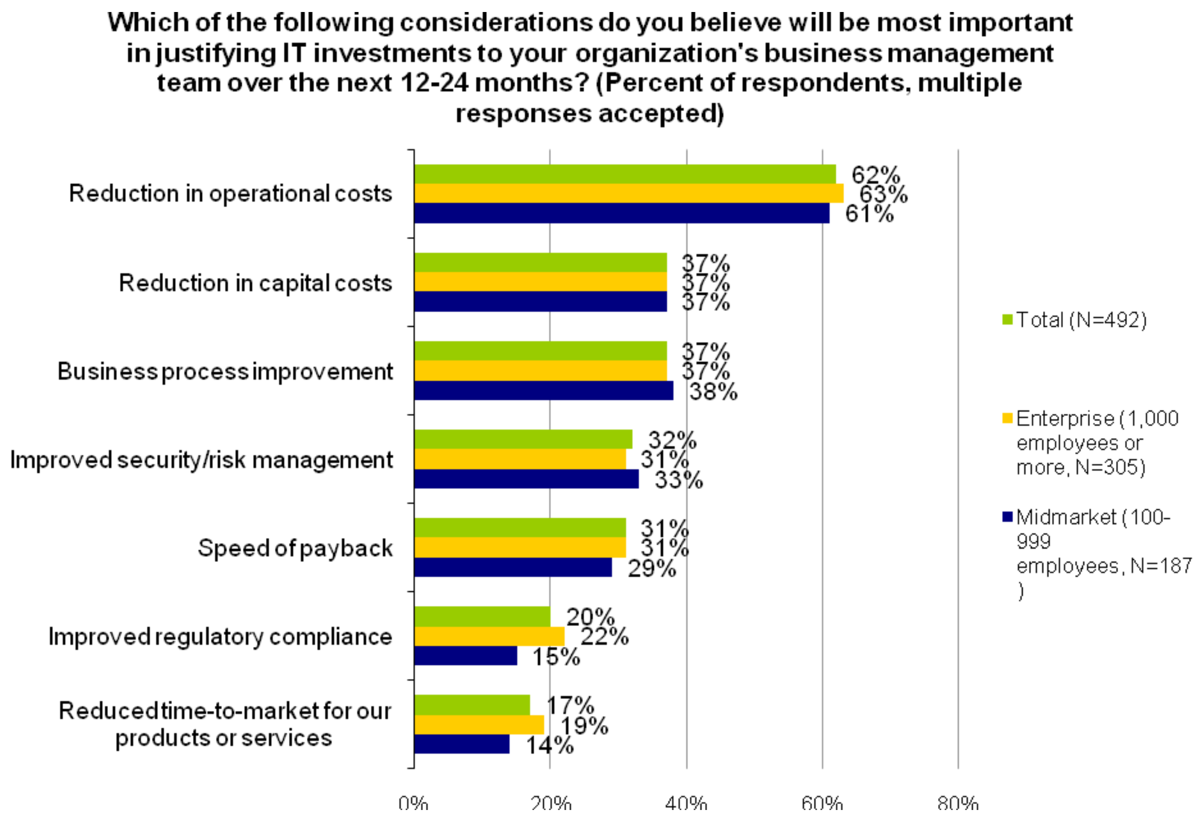

Source: ESG Research Report: 2009 Data Center Spending Intentions Survey, March 2009

While many other considerations were noteworthy, the emphasis on reducing operational costs is not only huge, but significantly ahead of any other consideration. And, again, because NetApp offers a choice of solid state solutions, users can adjust the balance and extent of their implementation(s) to suit their specific situation in terms of needs, budget, and comfort level. So, for instance, testing has shown that for a given IOPS target, a FC drive + PAM approach might reduce the drive count by 50% and give a 35% better response whereas a SATA drive + PAM choice might yield a similar reduction in the number of drives, but increase the overall capacity by 75%.

There is also a selection element when users want to consider the storage efficiency features that are available with different solid state storage implementations. Whether solid state is applied for caching purposes or as a storage tier, NetApp can help users increase the utilization of their storage and hence drive better value from their systems. Naturally, the varying use of solid state maps to different 'efficiency features'. All the implementations that NetApp is bringing to market (PAM cache, virtualized SSDs and native SSDs) can work with deduplication; virtualized and native SSDs can also benefit from thin provisioning, snapshots, and clones (the NetApp term for writeable snapshots); and, being an element within the disk enclosure itself, native SSDs will also be integrated in NetApp RAID-DP. Once again, solid state is not just one implementation—so how do you determine what's best?

Implementation Considerations

By now, it should be obvious that the first question a user should ask when considering solid state storage is not 'how fast?' nor 'how many IOs?' nor is it 'what does it cost?' Instead, the first question to ask when considering solid state is, 'what are you trying to achieve?' This might be followed by:

a. Do you have performance challenges? Increasing application expectations?
b. Is your server and application performance limited by storage latency and IO bottlenecks that compel you to 'throw' additional and expensive disk spindles or DRAM at the problem?
c. Do you know what your energy costs, usages, and objectives are?
d. Is your current storage infrastructure under-utilizing (a.k.a., wasting) resources?
e. How are you mitigating that situation now?
f. Are you aware of all the solid state options that you have?

From there, a plan can emerge. Once again, dealing with a vendor such as NetApp that a) 'gets it' and b) offers multiple solid state options is most likely to lead you to an optimal solution. Better yet, for its PAM family, NetApp offers a tool (Predictive Cache Statistics or PCS) that enables NetApp users to model and evaluate the potential benefits of adding more PAM cache to their system. PAM family modules can be set up in varying 'modes' to cache both data and metadata (its default) or only metadata. Since many IT managers are unlikely to precisely know the size of their most active data, or the applicability of SSD to particular applications, it can become a challenge to determine how effective SSDs will be in their environment *without* this sort of advice – and, after all, the last thing any user would want suggested by a vendor is to simply install something and hope for the best! The PCS tool from NetApp is designed to replace such a cavalier approach with pragmatic advice.

SSD is optimum mainly for persistent and active data with random intensive IO patterns where the aim is to guarantee that every read is fast. Such 'performance-centric' uses could include OLTP, database acceleration, HPC, and even demanding web services.

Cache (PAM II) is optimum mainly for random read intensive patterns where the aim is to improve the average response time across all the storage. Such 'cost sensitive' uses could include file services and collaborative tools such as MS Exchange or SharePoint, as well as home directories or VDI environments. As field experience increases, NetApp users have reported that caching is also effective for read-intensive databases and for technical applications including software development, electronic design automation, and product lifecycle management.

There are no absolute rules—as with the storage hierarchy in general, it is the specifics of a user situation (and many factors therein) that will determine what is possible and optimum. That said, there are very few places that could not benefit from some form of solid state to some extent. Over the next few years, some solid state storage will be a core part of almost any efficient data center. Natural inertia and skepticism in IT may extend this process, but the ultimate success of flash storage—not replacing HDDs, but complementing them—is a result of the simple logic of the storage hierarchy. NetApp makes the process easier for its users by already offering a 'hierarchy of solid state,' which is a natural and eventual direction for the whole market.

Conclusion & ESG View

From a market perspective, it is clear that solid state storage is moving from the edge to the mainstream. The basic promise is simple and undeniably appealing: it offers operational and business value through a combination of very high performance and reliability with low resource consumption (power, cooling, and space). But then comes the sticker-shock! And yet this is crazy in many respects—why be put off by a relatively higher cost for solid state compared to spinning disks? After all, there's no one who would look at the relatively higher cost of spinning disk compared to tape and suggest that's either a problem or a significant reason to store everything on tape! Once we acknowledge that solid state is simply part of the storage hierarchy, then figuring its value and optimum utilization becomes straightforward. Moreover, it is implicit that there should be a range of types of solid state, as there is with disks and indeed tape. It is surprising how many vendors seem unaware of, or silent about, this logic. NetApp is one of the few vendors taking a more thoughtful approach.

A 'more thoughtful approach' means that solid state is not one tier, nor is it one product; it is a choice of implementations and applications focused (for once and at last) on what really matters to active applications: IO, not storage. The value is already in the products; while this will improve, the fact is that many users could benefit from some level of implementation immediately. Waiting for the 'right moment' or the 'perfect price' is like trying to apply an absolute point to a relative target. Given that the economics of the storage hierarchy make solid state sensible for most users to some extent, the discussion today should be around 'how much' and 'how fast' rather than 'whether.' And, in the medium term, flash and SSDs will become significant within storage infrastructures—not as a percentage of capacity stored, but as a percentage of IO handled. Whether those IOs are of persistent data or of cached data, the benefits are similar: better performance, less disk reads, less physical spindles, better resource usage.

Turning to NetApp specifically, it is to be applauded for thinking broadly and not just following the herd to certify an SSD or two into its storage boxes. Of course, this should be good for NetApp's bottom line, for it definitely is offering additional end–user value by having choice and understanding, not just a hardware 'bolt in.' Its approach helps to move us toward an 'IO Hierarchy' because that is what solid state—in all its current and anticipated forms—is all about. When you move from performance as an end in itself to realize performance is simply a means to an end, you can begin to analyze solid state pragmatically. It is not voodoo, it is value. And it's not a thing, per se, but rather a whole family of technologies that constitute an extension of the storage hierarchy. Viewed in that manner, it is subject to all the same rules as any other storage offering: what do I need where and what can I afford to put where? The more options, the better—and, for now, NetApp has grasped the situation and offers at least as many, if not more, solid state options than its counterparts. From regular SSDs to large scale arrays to caching with PAM II, it's an impressive line-up that is worthy of closer inspection and consideration.

Enterprise Strategy Group

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com