



White Paper

EMC's Comprehensive Flash Strategy

By Mark Peters, Senior Analyst

July 2014

This ESG White Paper was commissioned by EMC and is distributed under license from ESG.

Contents

Flash Storage: Pragmatism versus Panacea	3
Flash Storage: Past and Present	3
History and Opportunity.....	4
Performance and Economics.....	5
'Horses for Courses'	5
EMC's Flash Implementations: A Strategy of Choice.....	6
Market Relevance and Value.....	9
The Bigger Truth	12

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188. Note that the first publication of this paper was March 2013; while some explanatory sections are very similar, this is a substantially updated version.

Flash Storage: Pragmatism versus Panacea

The flash storage market today feels somewhat reminiscent of a famous old movie line: “Flash, I love you! But we only have 14 hours to save the Earth!”¹ If you drink too much industry Kool-Aid, you could be forgiven for thinking flash storage is indeed the savior of the storage world. And there’s no doubt it can do wonderful things. But the truth is far more about pragmatism than panacea. *This is because flash is a technology, and not a market.* Indeed, flash is—more accurately—an umbrella technology type, within which there are numerous implementations.

Solid-state storage has been around in some form or another for over three decades, but only in the last few years has it become widely adopted, a change that is principally based on three things:

- **IT Needs:** The need for higher performance driven by such things as new applications, virtualization, the web, clouds, “consumerization”...and let’s not forget higher user expectations. There is an increasing “understanding” that performance will simply “be there.”
- **Media Attributes:** Solid-state technology itself has overall improved, with flash storage being nonvolatile, much less expensive than even a couple of years ago, and sporting much-improved manageability, longevity, and reliability (albeit the last two elements are invariably achieved via sophisticated management algorithms because the charge to higher capacity/lower cost raw media is not actually improving longevity).²
- **Deployment Options:** Coming to terms with flash media, the storage industry has developed methods to:
 - Judiciously leverage small amounts of this still-relatively-expensive resource to deliver great impact, both in terms of operational benefit and financial attraction.
 - Enable sufficient flash efficiency so that IT users can now run whole or multiple workloads in all-flash configurations while still enjoying economics that are comparable to hybrid systems.

And to be crystal clear, flash is no longer *only* about performance, albeit that performance can be necessary and useful. In the same way that we accept a range of disk drive and other storage types, we should understand and embrace that we will have a hierarchy of flash (today) and solid-state storage technologies (tomorrow).

[EMC](#) kick-started the “new era” of flash when it introduced Enterprise Flash Drives (EFDs) on its Symmetrix arrays in early 2008. Now every storage vendor offers flash in some manner and a whole host of start-ups have been spawned, with many already snapped up by “the big boys.” EMC itself has been busy in flash, seeming to double-down on its investments in the segment regularly, such that it has a plethora of flash options that is now available, announced, or—with its latest DSSD news—acquired and intimated. Is it just being a corporate magpie, or is there strategy and intent in EMC’s “embarrassment of flash riches”? After all, some users might feel that just one or two flash options should be enough. How fast is fast enough!?

The top-line answer is that EMC is demonstrating, once again, its core competency for dramatic—and necessary—*change*. It has appreciated the broad, long-term strategic importance of flash storage and is moving rapidly to both maximize and optimize its use across the EMC portfolio. Before looking at the EMC line-up, and evaluating its applicability and logic, let’s take a quick “Cliff Notes” look at the world of flash.

Flash Storage: Past and Present

It’s impossible to conduct a storage conversation these days without talking about “solid-state” or “flash storage.” Along with “virtualization,” “cloud,” and “software-defined,” it’s one of those de-rigueur phrases. And, like those other two, flash is suffering from an over-exposure that—even now, nearly seven years into the “new NAND era”—can assume familiarity and understanding where it does not exist.

¹ Source: A quote from the script of the movie *Flash Gordon* (1980).

² Notably, the general and overall “reliability” of NAND flash was initially considered an Achilles Heel for the technology, but is now a de facto advantage with most anecdotal user—and private vendor—data reporting levels of reliability that have surpassed initial expectations.

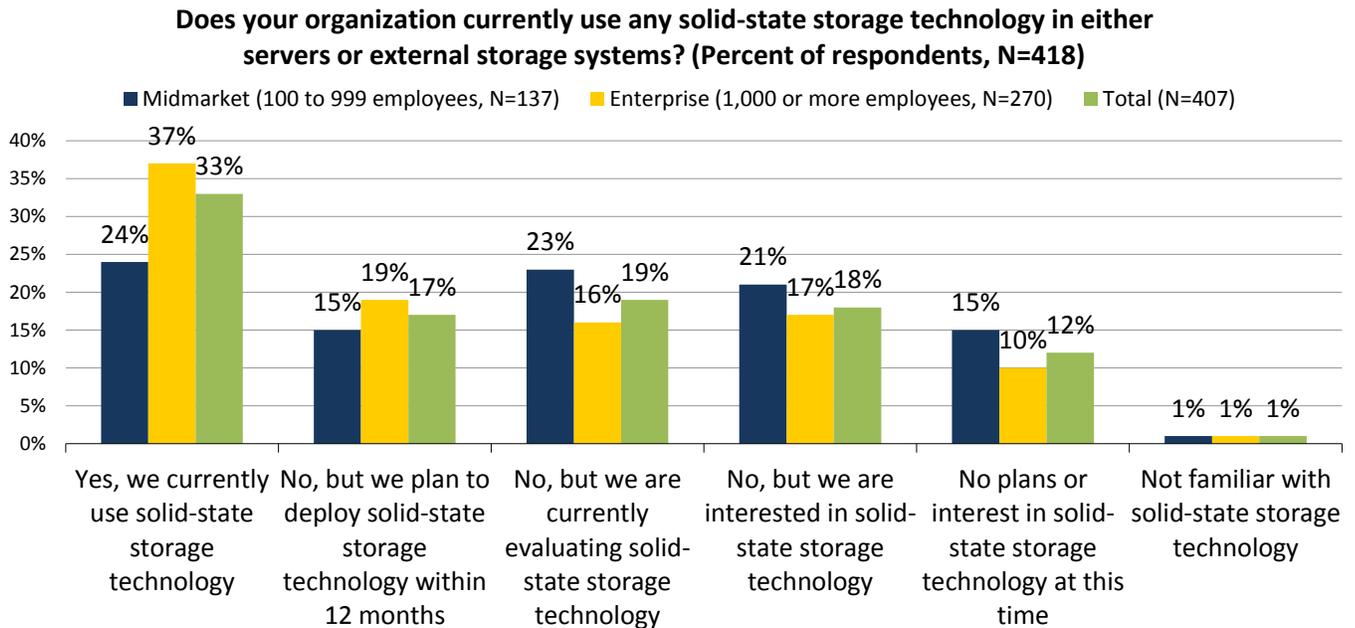
History and Opportunity

Solid-state storage is far from new. It first appeared as a dynamic random access memory-based (DRAM) device back in the late 1970s and has been a popular, but limited, niche product ever since. Flash memory (NAND) has changed all that by presenting the market with an option that is nonvolatile in nature, not stratospheric in price, and which innovative vendors are packaging cleverly in order to drive function and economic value. If hype alone determined market adoption of any given technology, then solid-state storage would already have taken over much of the storage world! Certainly, the attractions of flash are pretty apparent: performance (a “means” as well as an “end”), compactness, reliability, and low power and space consumption.

Yet, despite a lot of mathematics and marketing to the contrary, there remains a significant general raw price differential between flash and spinning disks.³ As users grapple to digest the opportunity, vendors are laying out a wide assortment of solid-state storage options. These can be segmented into three main infrastructural implementations: in the server (as direct-attached storage (DAS) or cache), in the storage subsystem (controller or array), or as a standalone appliance/all-flash array (AFA). There are also two main usage options: persistent storage (a tier) or temporarily stored copies of data (a cache). ESG research confirms the popularity—both existing and expected—of flash storage (see Figure 1).⁴

Since that formal research was conducted, anecdotal reports from the user community strongly support the notion of an even higher adoption rate and storage vendors regularly confirm that—aside from the dramatic uptick in AFA sales—the vast majority of their “regular” contemporary storage system shipments now contain at least some amount of flash.⁵

Figure 1. Solid-state Storage Adoption Trends, Overall, Midmarket and Enterprise



Source: Enterprise Strategy Group, 2014

³ Cases can be constructed for flash to be a better price/GB than disk, but invariably rely on using data compaction and a comparison to only “high-performance disk.” While not unfair, these cases are handpicked and often obfuscate the point by unnecessarily honing in on \$/GB.

⁴ Source: ESG Research Report, [2012 Storage Market Survey](#), November 2012.

⁵ For 2013 EMC, for example, shipped flash in 76% of its VMAX and 71% of its VNX devices: Looking at the data in Figure 1, and using EMC as a market proxy, this would tie well to the notion that the research was accurate and that those IT users who reported planning or evaluating flash in late 2012 have indeed moved to adopt it.

Performance and Economics

For decades, traditional spinning hard disk drives (HDDs) have had to deliver both performance and capacity (frankly without excelling at either, but being the least-worst option available!). While there have been noticeable advances in HDD capacities, the actual performance has essentially not improved, and this is at the same time that processing abilities and demands have sky-rocketed. The fact that flash can add performance across the storage environment is very welcome—but *it is only gaining ground because it can do so while also making economic sense*. There are two reasons for this:

1. Today's HDD technologies *can* deliver performance—but typically in a manner that leads users to do economically illogical things. Short stroking, for instance, means users buy extra HDD spindles for the I/O they gain, even though they don't need the capacity. That's a waste of resources and an effective increase in the \$/GB of all the storage in a system.
2. All storage decisions come down to a financial decision eventually. The storage hierarchy—and the choices it drives—only exists because storage isn't free. If storage were free (or at least all the same price) then logically we would put all persistent data into memory. There are of course important niche areas of the market where performance is paramount and where the cost doesn't matter (electronic trading being a popular example), but for the vast majority of users, there is a cost-benefit analysis to be made for flash to be adopted. *Perhaps counterintuitively, it is the economic advantage that flash represents (rather than its performance ability per se, but what impact that performance has on the overall economics of a user's storage environment) that is increasingly driving its adoption and success.*

It's obvious what high-performance storage can do for applications that drive high I/O to handle extremely frequent transactions or that are ultra-sensitive to latency. However, what's also invariably true is that a judicious amount of high I/O and performance capacity (whether as a small percentage within a *hybrid system* or as a 100% all-flash array, albeit this is often then part of a *hybrid infrastructure*) can do wonders to drive better storage economics. *Put simply, extremely high-performance storage—invariably flash of some sort—can be used to manage costs!* Of course, even economic advantage can be delivered in a number of ways. Desired outcomes are not only about being *cheaper* (that's just on the “cost” side, but can of course include myriad OpEx costs as well as CapEx), but are also about delivering *business value* (which is on the “benefit” side and includes such things as higher margins, faster time to market for new products, improved customer service, and so on).

'Horses for Courses'

It seems clear that solid-state of some sort *is* the future for all sorts of general purpose storage...and in centuries to come, we will likely look back at the first few decades of IT and smirk at the notion of placing our most precious data on spinning rusty platters and low-flying magnetic heads. Economics matter; increasingly the industry is talking about \$/IO or TCO/I/O as well as just \$/GB, but IT remains a pragmatic industry where cost matters greatly, especially as the demand for its services seems elastic.

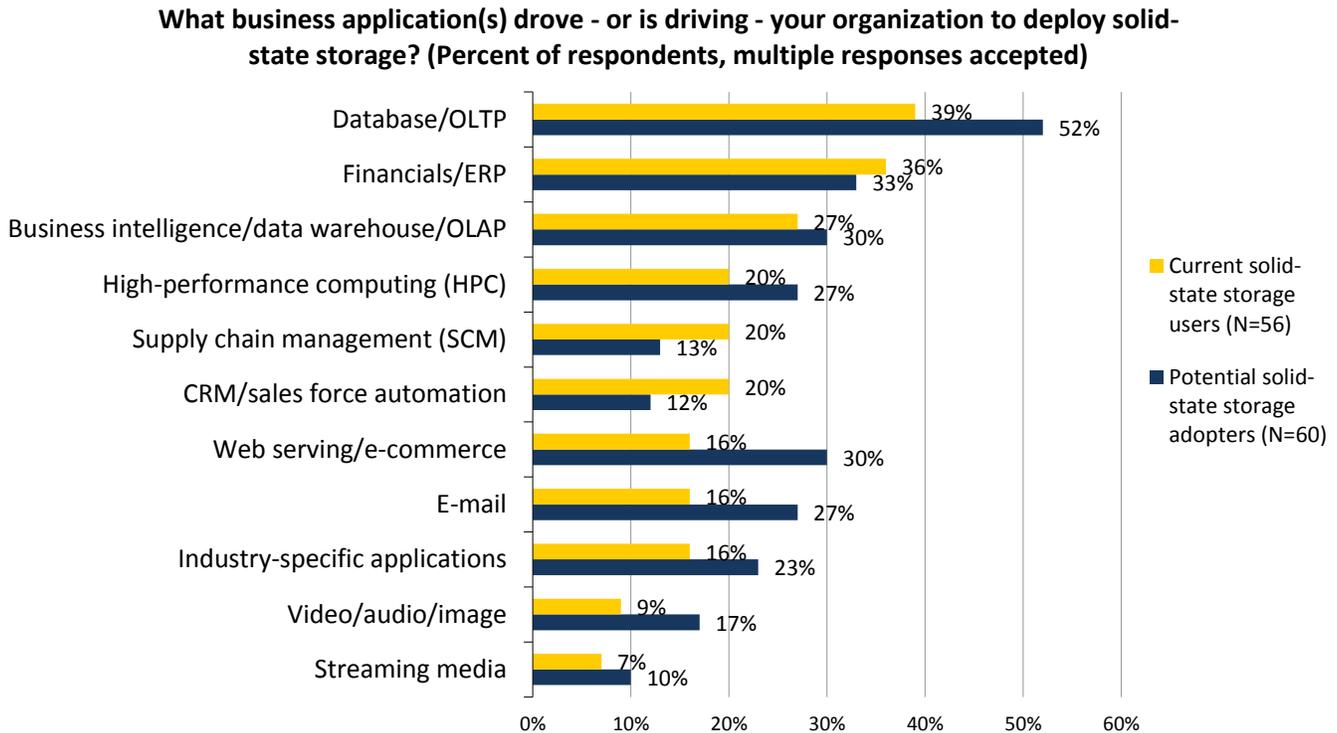
Thus, for the foreseeable future, we are business-constrained to work with a storage hierarchy and the careful application of flash wherever it can help. Again, this is foundational for EMC's broad flash portfolio. Flash may not be taking over the *storage capacity* world any time soon, but it can categorically be used to serve and improve an ever-increasing percentage of the *storage activity* world. That, in turn, demands that flash be implemented in numerous ways, so that it can serve different uses, applications, and requirements.

In this light, EMC's wide, and indeed increasing, range of flash products is not simply about backing all the horses in a race; more accurately, it is a recognition—to extend the equine analogy—that there are many workloads that demand different types of horses. And there are many types of “races.” You would no more expect a shire horse to run in a Derby than expect a thoroughbred racer to pull the beer cart. More pointedly, even purebred race horses are bred and trained for different distances and styles—sprints, endurance, and flat or fences. And even the more specialist “high-performance” requirements can vary: Dressage and ceremonial work are not the same. These examples explain the well-known term “horses for courses.”

EMC's Flash Implementations: A Strategy of Choice

The Value of Options: Moving from horses to EMC's flash portfolio, the essence is nonetheless similar; different horses (products) for different courses (applications and workloads). Each workload requires different things from flash and, therefore, is invariably best served by different implementations of flash. For years, IT users have bought and planned flash storage to serve a broad range of applications, as shown by the ESG research data in Figure 2.6

Figure 2. Specific Applications that Drove—or are Driving—Solid-state Storage Deployments



Source: Enterprise Strategy Group, 2014.

This increasing range of applications is happening at a time when server virtualization is endemic: This inherently drives up the performance needs from each storage device because of both the consolidation and the broader range of applications that are likely to be relying on any given storage system. Put simply, today's flash—of whatever implementation—is less about being a “point” solution and increasingly about having horizontal flexibility. There are still different courses and race types (so having multiple horses is still the answer), but we are looking for more cross-over training and operational flexibility.

Indeed, having spent decades finding ways to mitigate the rotational latency and other constraints of HDDs, the industry has moved from optimizing for spinning disk to optimizing for flash, and that is allowing us to do myriad new things (think, for example, of the ability to create thousands of snapshots without degrading performance) that have released application vendors and users to discover what is possible when the prior constraints are withdrawn. This is not to suggest that flash does not bring its own challenges, but they are generally more predictable and manageable, and can therefore be addressed via system architectures that account for things like garbage collection, write wear, write amplification, and memory management.

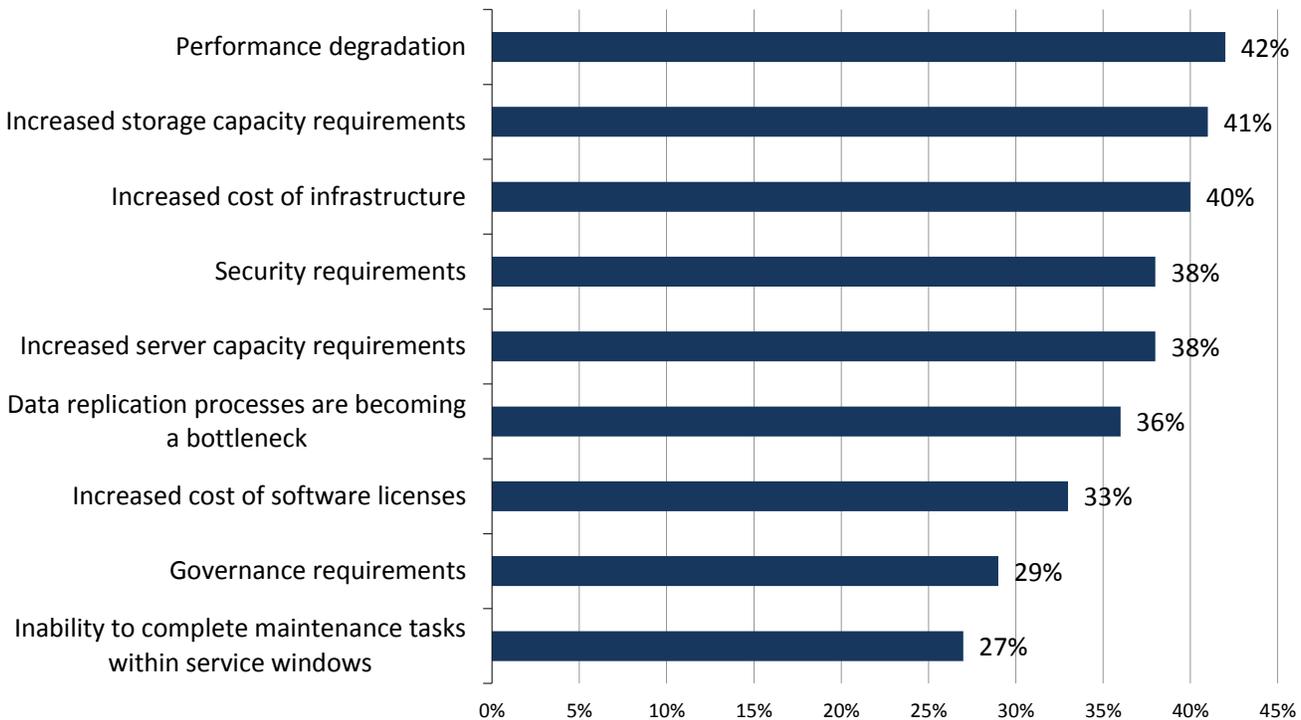
Nonetheless, whether applications run on a physical or virtual machine and whether they are in a traditional data center or a (hybrid) cloud operation, at a 50,000-foot level, those with “high-intensity I/O” needs broadly fall into one of two camps: *low latency* needs, such as VDI, e-commerce, in-memory compute, and real-time analytics, or *acceleration* needs, such as databases, OLTP, and decision support systems. Databases are an interesting subset

⁶ Source: ESG Research Report, [Solid-state Storage Market Trends](#), November 2011.

because just about all businesses operate databases of some sorts, and—as the recent ESG research shown in Figure 3 demonstrates—performance remains the most-cited challenge for database users.⁷

Figure 3. Challenges Specific to Database Size and/or Growth

For organizations experiencing difficulties related to database size and growth, what are the specific challenges faced with respect to database size and/or growth? (Percent of respondents, N=119, multiple responses accepted)



Source: Enterprise Strategy Group, 2014.

Moreover, the same research validated that the average production database for the vast majority (84%) of respondents is under 10TB and even the largest production database is 25TB or less for 87% of respondents. The reason these numbers are significant is that, at these levels of capacity requirements, and with the functional sophistication of current AFAs (such as EMC's XtremIO), there is every good reason—and affordability—to place entire databases on flash storage.

Latency and acceleration are the two broad attributes of flash, within which the current EMC product line-up offers a breadth of variation in terms of function, price, raw speed, and so on in order to meet a huge range of buyer and workload needs. The main categories in today's portfolio are the hybrid versions of the leading storage platforms and the XtremIO all-flash array, but there is a comprehensive choice both within categories and surrounding them.⁸ To return to the earlier analogy, if the hybrid systems are viewed as workhorses, then XtremIO systems represent race horses: indeed, like truly exceptional horses, lifetime success is measured not just by always being the fastest or by always running the farthest, but by being able to handle multiple and varying courses and conditions (workloads) as they arise. Increasingly, of course, users are not choosing between flash deployments as exclusive alternatives, but rather are choosing combinations among them to inclusively address their many, varying needs.

⁷ Source: ESG Research Report, *Enterprise Database Trends*, to be published July 2014.

⁸ This paper is not intended to be a substitute for EMC datasheets; it merely gives some product logic, headlines, and brief explanations.

- **All-flash Arrays⁹:** With enterprise data services such as snapshots, replication, and encryption in addition to such efficiencies as deduplication and compression, an increasing wave of users are now choosing to consolidate their production applications onto all-flash platforms like EMC's XtremIO. The XtremIO product offers persistent storage, a tightly coupled linear scale-out ability (for consistency via the shared memory space control), sub-one millisecond latency, and a "zero planning/tuning" content-based placement (to preclude hotspots and enable deduplication). The scale and mission-criticality of AFAs such as XtremIO mandates an intense engineering focus to optimally manage data layout, placement, and addressing to eliminate all unnecessary back-end I/Os. The EMC XtremIO answer is a software approach that runs exclusively in-memory, which means that all data management—such as deduplication, compression,¹⁰ and XtremIO's advanced, thinly-provisioned and writable snaps—is handled inline in DRAM before anything is written to flash.¹¹ Specific capabilities such as "Dual-Stage Metadata" and "XDP" enhance system efficiency by, respectively, eliminating the overhead of system-level garbage collection and minimizing the capacity that is needed for parity data protection.
 - **Typical Workloads:** Applications requiring highly consistent sub-one ms performance but still with random I/O that are too large for server-based flash or that don't respond well to caching and tiering algorithms are suited for an all-flash array. These might include VDI, virtual servers, and databases (throughout their life cycle of needs, including test/dev, production, analytics, and backup).
- **Hybrid Arrays:** Not all workloads need the attributes of all flash all the time; in such situations and for IT users with hundreds of TB to PBs of storage capacity, a little flash can be made to go a long way and the economics of hybrid arrays become very favorable. EMC's traditional headline storage systems (VMAX, VNX, VNXe, and Isilon) can all be deployed with variable amounts and percentages of flash, and all feature automated data placement (i.e., hot data goes on flash and cold data on HDD) to provide a balance of (usually not the absolute highest) performance and (often the lowest) cost. These are typically best for users' persistent and predictable data that represents much of the basic daily work in most organizations.
 - **Typical Workloads:** Most suited for larger data sets and mixed, variable or consolidated workloads with some tolerance for occasional latency (still good but in the 1-2ms range). These might include data warehousing, OLTP, and e-mail, but the relative applicability for any workload will always depend on individual user situations.
- **Server Flash as Local Storage:** EMC XtremSF is a high-capacity PCIe-attached, media card (either SLC or eMLC, depending on the workloads) deployed in the server as local storage for extreme application performance acceleration: This improves the performance of application reads and writes by dramatically reducing latency and speeding throughput.
 - **Typical Workloads:** EMC XtremSF is suited to high-transactional and high-performance workloads often associated with web 2.0 applications, smaller virtual desktop infrastructure (VDI) environments (say =/< 100 desktops), high-performance computing (HPC), and high-performance trading applications. It can also be used to accelerate analytics, reporting, data modeling, indexes, database dumps, batch processing, background tasks, and other temporary workloads.
- **Server Flash as Cache:** For mission-critical workloads in which data must be protected by the array, XtremSF can be coupled with intelligent server flash caching software—this is XtremCache. While accelerating reads, XtremCache leverages a write-through algorithm to ensure that newly written data persists to the networked storage array for persistent high availability, integrity, reliability, and DR.

⁹ While XtremIO is clearly EMC's leading AFA, July 2014 also saw the arrival of VNX-F, which is an all-flash VNX configuration. It is essentially "flash by the pound" for users who want performance but do not require any data services, and it appears to mainly be EMC ensuring that it can offer every flash variant, and compete across the "portfolio board."

¹⁰ Announced July 2014.

¹¹ This approach minimizes and pre-handles all activity that results in writing to disk, such as buffer space for batch processing, metadata de-staging, or post-process data services.

- **Typical Workloads:** XtremCache is optimized for web applications, online transaction processing (OLTP), customer relationship management (CRM) and enterprise resource planning (ERP) databases, e-mail applications, and other read-intensive workloads with small working sets.
- **“Rack Scale Flash”:** This is the descriptor that EMC is currently applying to what will result from its acquisition in May 2014 of a company call DSSD. As the scale of in-memory workloads and high-performance analytics grows, EMC is taking the next step for solid-state (flash is only one element here) by making a very large-scale flash-based tool that exists either as DAS or an appliance and operates essentially as memory extension; to achieve this, it combines extreme density and massively parallel connectivity (perhaps 8,000 or 16,000 independent connections).¹²
 - **Typical Workloads:** The eventual product here is expected to be targeted at in-memory database processing, and real-time/high-performance analytics; thus being for things like MongoDB or running a high number of SAP HANA instances.

Beyond the hardware, a growing part of EMC’s differentiation and strategy (and likely a glimpse at the future in our newly software-defined world) is software. This is not only the caching algorithms and management protocols for all the hardware mentioned, but also EMC’s Fully Automated Storage Tiering (FAST) that works on EMC’s hybrid arrays to get the most active data at any time to the most suitable, available, and higher performing storage. Applications with mixed workloads and changing data “temperatures” are typically a good fit for FAST as part of a hybrid storage tiering strategy. Looking ahead, it seems likely that EMC will extend FAST to move data between arrays, not just between tiers within the array. Beyond this, various ingredients for a much wider, more fluid, granular, and integrated abstraction can already be seen in such tools as ScaleIO (another recent EMC acquisition that virtualizes and pools server-based storage) and ViPR, both of which are part of EMC’s software-defined data center strategy.

Market Relevance and Value

This raises an important point: It is all too easy to get carried away on the joint waves of performance and—yes—coolness of flash. But users will balance their solid-state evaluation and adoption with two large elements of pragmatism: while things like TCO absolutely matter in selecting storage vendors and solutions, an easy-to-manage and low-complexity solution (the sort of thing that XtremIO, or FAST-based EMC hybrid arrays promise) is also important, as the ESG research in Table 1 shows.¹³

Table 1. Important Criteria When Selecting Storage Vendors/Solutions, by Company Size

In general, what would you consider to be the most important criteria to your organization when it comes to selecting a storage vendor/solution?		
	Midmarket (100 to 999 employees, N=141)	Enterprise (1,000 or more employees, N=277)
Ease of management	56%	42%
Ease of implementation	50%	36%
Ability to leverage existing staff skills	24%	33%

Source: Enterprise Strategy Group, 2014.

The Need for Options: If the market had already chosen a clear winning solid-state implementation style (which is itself unlikely because of the differing needs that are being served), then EMC could perhaps make do with fewer flash offerings. However, both current flash deployments *and* those that users expect to make demonstrate that

¹² EMC has not publically or officially said much regarding DSSD at the time this paper is written; the notes here are sourced from discussions at EMC World, the Virtual Geek blog posted in May 2014 (written by Chad Sakic of EMC), and reported comments by Dan Cobb, the CTO of EMC’s flash business, in a ComputerWeekly.com article posted in June 2014.

¹³ Source: ESG Research Report, [2012 Storage Market Survey](#), November 2012.

the market has an appetite for a broad range of flash implementations. Again, given the equally wide range of users, needs, budgets, and workloads, this should not be a surprise. As a large, broad, and leading storage vendor, EMC would be “cutting off its nose to spite its face” if it were not able to meet as many market needs as possible.

There have, nonetheless, been changes since flash first came to the fore. The use cases for flash have changed over its short time in the market. Early workloads were often persistent and focused on things such as HPC, database acceleration, decision support for data mining, and metadata. These have been joined by less persistent workloads that are enabled by auto-tiering and caching tools with more horizontal applications (such as server virtualization, VDI, and consolidation). Changes naturally go hand in hand with changing types of flash: For instance, less expensive flash media (the cost is leveraged further down by data reduction technologies) that is better optimized is helping to rapidly grow the use of all-flash arrays (AFAs) as workload consolidation platforms, especially in organizations that have overall lower storage capacity needs. Inline deduplication and compression technologies have brought AFA configurations into the appropriate \$/GB range and are sometimes even lower cost than those of hybrid arrays.

Mentioning \$/GB again—after having previously highlighted its lack of sophistication and sufficiency as a decision metric for flash purchases—is entirely deliberate. *Cost still matters*: Whether as \$/GB or \$/IOP or whatever, cost remains the most-cited business initiative that drives technology expenditures, according to the respondents in ESG’s latest Spending Intentions research (see Figure 4).¹⁴

Figure 4. 2014 Business Initiatives with the Greatest Impact on IT Spending Decisions

Which of the following business initiatives do you believe will drive the most technology spending in your organization over the next 12 months? (Percent of respondents, N=562, five responses accepted)



Source: Enterprise Strategy Group, 2014

¹⁴ Source: ESG Research Report, [2014 IT Spending Intentions Survey](#), February 2014.

This insight about the importance of cost reduction *matters very much for the current and future deployment of flash technologies*. The simple fact is that intelligent deployments of flash *will just about always save money for IT users*. This is not some loose assertion that is full of “apples versus oranges” caveats. It is simply that a well-implemented, judicious amount of flash (whether a simple card, a hybrid, or an all-flash array) is overall a less expensive way to deal with demanding I/O needs. It therefore provides a better ROI for IT departments; and this matters, too, since the cost reduction need that is highlighted in Figure 4 is not about slashing budgets: The same 2014 spending intentions research found that the most commonly leveraged measure for cost containment involves purchasing new technologies with a better ROI.¹⁵ Furthermore, ROI is the most-cited consideration for IT organizations to justify their IT investments to their business management team, followed very closely by reducing operational expenditures and improving business processes...all three clearly being considerations that good flash deployments will impact positively.¹⁶

And this race is still being run. Looking ahead, the most successful vendors seem certain to be those that offer—in addition to a range of solutions—some share-ability of flash resources; a hybrid approach, which means in terms of infrastructure as much as in a single box solution; and comprehensive management, optimization, and orchestration software. EMC's extensive flash portfolio approach seems to have the requisite abilities—shipping, announced, and nascent—on all these fronts.

¹⁵ Ibid. Renegotiating contracts was cited as often as purchasing new technologies with improved ROI (36%) as a cost containment strategy.

¹⁶ Ibid.

The Bigger Truth

Winston Churchill once said, “Democracy is the worst form of government there is...except for all the others.” And that’s how disk storage has been, too!

Flash is now rapidly and violently disrupting the disk status quo, and not just, or even mainly, because it is fast. Yes, the early integration of SSDs into a storage strategy was all about accelerating specific workloads, followed by getting a better overall balance of price/performance in external storage arrays. However, as flash is becoming more pervasive, optimized, and affordable, new applications and new storage software stacks are also emerging that are built “ground-up” for solid-state. Thus today, flash storage has value in a number of implementations within and across the storage infrastructure—be that in storage systems themselves, in servers, or in the network. Since the cost of flash still makes it uneconomical for most users to switch over completely (although even that is beginning to change in some places, and the new arrival of capabilities like data at rest encryption on XtremIO shines a light on the future), the simple fact is that using a small percentage of flash in almost any location in the storage hierarchy, and using it in combination with some intelligent management software (optimization code in an AFA, and caching or tiering in hybrid arrays) will almost always yield both performance and financial improvements.

Since economics underlies all storage decisions, it follows logically that most users will continue to maintain a hierarchy of storage across their infrastructures as a whole. Moreover, any single implementation of flash of just one style and type in any given user would almost certainly be sub-optimal compared with a multi-flash implementation at the same user. This is not to say that single-flash implementations cannot yield improvements...it’s simply that multi-flash models will invariably be able to deliver *even better* results. That, in a nutshell, is the IT and user-value logic behind EMC’s comprehensive flash portfolio.

Obviously for EMC, the business motivation is to be as much of a “one-stop flash shop” as it can, in order to retain its storage leadership position. Even if occasionally there’s a small positioning issue as to what is the best solution for a given user’s situation, the product overlap is far preferable to a product gap, or having just a single product (in each case, the old adage about “when you only have a hammer, everything looks like a nail” applies to some degree). To put this in other words: If you were a new storage vendor with the same options, skills, and resources as EMC, you would be doing exactly what it is doing.

From the earliest days of its involvement with flash, EMC has made it clear that it is fully committed for the long haul. This is a matter of corporate faith; it is not about the next flash widget or a single flash product (indeed, in the coming years, it may be about some form of solid-state other than flash), but is instead about understanding that having multiple offerings makes sense when both the capability *and* applicability of flash covers such a broad spectrum. EMC has always shown a willingness to, and aptitude for, change; even if flash means it will have to “eat its own disk children” to some degree, EMC understands that it’s necessary in order to retain its market leadership. Indeed, something very interesting happened when EMC initially announced its Xtrem Family: While almost all the focus was on the *product* that would initially follow, what was actually far more significant then (and now) was that EMC had established a focused flash *program* and division. In a company that loves its acronyms, that move told us everything we needed to know: “FLASH” for EMC is not about any particular product; instead, it seems to stand for its intent to Fully Leverage A Silicon Hierarchy!



Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com