CLABBY ANALYTICS

Research Report

EMC's Drive into the Flash Storage and Flash-as-Memory Marketplaces

Executive Summary

We look at "Flash" technology (a type of erasable non-volatile computer storage medium) as "fast storage" and "slow memory". What we mean by this is that Flash is being deployed in solid state disk (SSD) storage to improve reliability and deliver fast input/output per second (IOPS) performance; and it is being deployed as a new "memory tier" that places data close to a microprocessor to speed the processing of in-memory and Big Data workloads.

We now see a very compelling case to adopt Flash-based SSDs in place of hard disk drives (HDDs) based on five factors: 1) cost reductions; 2) performance advantages; 3) energy savings; 4) reliability; and, 5) better utilization.

We also see a very compelling case to adopt Flash memory as a new memory tier that can provide processors with instantaneous, real-time access to multiple terabytes of data contained in Flash memory. We described one such system design where Flash sits on a memory channel within an x86-based system design and acts as lower cost main memory in this report. And we showed other examples of Flash acting as a secondary memory tier in this report.

In this *Research Report*, *Clabby Analytics* examines EMC's Flash storage and Flash memory strategies. And what we find is that: 1) EMC has been investing in Flash research since 2005; 2) the company brought its first Flash storage products to market in 2008; 3) the company acquired XtremIO to accelerate its deployment of Flash-based arrays (and, according to *The 451 Group* and International Data Corp (*IDC*), EMC now holds the number 1 market position in all-Flash arrays in 2014); and, 4) the company acquired DSSD to accelerate its move into the evolving Flash-as-a-memory-tier marketplace.

> ***From our perspective, EMC is very well positioned to maintain a large share of the Flash storage market – while, at the same time, being well positioned to become a key supplier of Flash-as-memory products.***

The Flash Storage Marketplace

Over the past decade, processor, memory, networking and bus speeds have all increased dramatically. Thanks to multi-core processors, CPU performance has increased by a factor of 10. Memory speed has also increased by a factor of 9. Networking speed has improved by an amazing 100X. And bus speeds have increased by 20X. *In stark contrast, hard disk drive (HDD) read/write performance has only marginally increased, improving 1.2X over the past ten years*. So clearly, storage performance is begging for improvement. High performance Flash-based solid state drives (SSDs) represent that performance improvement.

Performance advantages, however, are not the only factors driving the computing industry to migrate from HDD to SSD. We see the following five factors that play a role in this migration:

1. *Decreasing cost of SSD storage* – the cost per gigabyte for solid SSDs has fallen steadily since 2013. In some cases Flash-based arrays can actually cost less than equivalently configured HDD arrays;
2. Faster input/output per second (IOPs) *performance*;
3. *Significant savings in energy use as well as cooling*;
4. *Greater reliability*; and,
5. *Better Utilization*.

Each of these factors deserves closer scrutiny.

Decreasing Cost

Since 2013, the cost of Flash-based SSD storage has dropped from about $2.00/gigabyte to about 69¢/per gigabyte as of last quarter – according to recent data published by IDC. IDC also projects that SSD cost per gigabyte will decrease to about 18¢/per gigabyte by 2017. Meanwhile, the HDD cost per gigabyte should decrease to around 3¢/per gigabyte in the same timeframe.

Many information technology (IT) buyers will note that, from a cost of acquisition point-of-view, in 2017 Flash-based SSD will still be about *six times* more expensive than mechanical drives on a dollar per gigabyte basis. Other factors, however, serve to reduce the cost of SSDs – in some cases making SSDs less expensive than HDD solutions. For instance:

- SSDs use significantly less power and require significantly less cooling than their HDD counterparts (so power/cooling should be part of any SSD vs. HDD comparison);
- HDDs are rarely used to full capacity whereas SSDs can be used to full capacity. As a result, HDD storage looks less expensive, but a closer look at the amount of data actually stored may reveal that as little as half of the available storage is actually being used. Or, another way to phrase this is that twice as many HDDs may be required to store the same amount of data that resides on an SSD; and,
- SSDs are far less expensive on a dollar per IOPS (input output per second) perspective (the next subsection explains why).

Performance

SSD technology very significantly outperforms slower mechanical drives (as shown in storage research studies found here, here, and here). These studies examine IOPS, throughput and response times – and SSDs win every time in these categories (one study showed an SSD performing at 44,000 IOPS – while a hard disk provided only 180 IOPS). The first study by Calypso is the most thorough – it shows a 15000 evolution per minute hard drive providing 401 IOPS (input/output per second), while an enterprise-class server using a PCI-e connected solid state card performed at 684,284 IOPS. Your mileage will vary based upon the hard drives in use, their speed and age, and the type and make of SSD that is chosen. But rest assured, SSDs, by virtue of being solid state electronics, always deliver data faster than mechanical hard drives.

It should also be noted that SSDs offer specific performance benefits to applications that exploit IOPS speed. Tech Republic hosts a study provided by LSI that examines the characteristics of solid state drives and describes the types of applications best served by SSDs. LSI states that SDDs "*have a faster start up, extremely low latency, near zero noise pollution, greater temperature range, higher reliability, and faster seek times*." Further, LSI states that "*SSDs provide excellent performance benefits for applications comprised primarily of random access patterns, such as online transaction processing, databases, and web servers*."

Energy/Cooling Savings
According to research provided by iSuppli, if SSD's had been used in datacenters between 2008 and 2013, datacenters could have reduced their cumulative electricity consumption by 166,643 Megawatt Hours (see this graphic). Much of this power saving comes from using SSDs that consume far less power than traditional HDDs (typically about half the power); but some of this saving also comes from energy saved by not having to cool large populations of hard disks.

Greater Reliability
There is an industry belief that SSDs are more reliable than HDDs because SSDs have no moving parts that could fail. But counter to this industry belief, SSDs do have parts that can fail such as transistors, capacitors and firmware.

It is hard to argue the case that all SSDs are more reliable than HDDs – because a lot of factors influence reliability. Enterprise Storage Forum points out, the reliability of a given HDD or SSD depends on several factors including drive "usage, physical environment, application IO, vendor, mean time before failure (MTBF)". But we can argue that some SSDs are more reliable than some HDDs due to fewer moving parts; better tolerance of hot and cold environments; and better design/manufacturing/quality control processes (vendor dependent).

Better Utilization
Most IT managers are reluctant to fully populate a hard disk with data – and, accordingly, this means that not every gigabyte purchased is being used. For those who are running mission critical data, perhaps only the first few tracks on each drive are used ─ and thus, perhaps only a third or half of the available mechanical disk storage is used. Users of solid state drives, on the other hand, are quite willing to use entire drives for storage because they don't have to worry about having to wait for the disk to spin to the location that they need it to in order to retrieve mission-critical data rapidly (SSDs don't spin) – and because solid state drives are more reliable (fewer moving parts). Also consider this: SSDs use only a fraction of the power used by mechanical drives – so power usage should be part of any hard drive/SSD price comparison.

IT executives are now buying SSD technologies to serve high throughput/low latency applications – and they are willing to pay a little bit more for SSD arrays. As these executives deploy SSD arrays, they are seeing significantly less energy consumed – and the resulting savings often close the price gap differential between SSD and HDD technologies. Further, these executives are finding that SSDs are more reliable than HDDs (because HDDs are mechanical and have moving parts – SSDs do not); and they are also finding that they can get better utilization out of their SSDs (administrators tend not to fully populate HDDs with data [which, incidentally, means that unused sectors are being paid for – but not used] – whereas SSDs can be reliably fully populated). Taking into account all of these factors, the total cost of ownership for SSD-based or hybrid SSD/HDD arrays is often lower than HDD only solutions.

Note that in the first quarter of 2014, EMC sold more than 17 petabytes of Flash capacity (up over 70% over the first quarter of 2013) – proof that EMC customers are embracing solid state storage architecture.

The Flash-as-Memory Marketplace
The situation in computer systems design has changed over the past few years. From a systems performance perspective systems makers used to rely on a process described as Moore's Law to increase computing speed by placing more transistors on a chip using a shrinking process. But Moore's Law has reached physical limitation as most commercial processors cannot achieve processing speed that goes much beyond 5 GHz. As a result, systems makers have turned to other computer systems elements such as improvements in the speed of the input/output subsystem and improvements in memory channel speed to increase overall systems performance.

Accelerating I/O: The Impact of CAPI and QPI

A well-balanced, well-tuned system does not waste clock cycles. So, to ensure that processors are kept busy, the I/O and memory subsystems need to feed the processors a steady stream of data to be processed. But traditional approaches to I/O communications are fraught with communications and management overhead that impedes the flow of data to its destination. For instance, using the tradition I/O flow model, device drivers are called, copy or pin source data is derived, memory mapped I/O needs to take place, acceleration needs to take place, polls and interrupts need to complete, copy or unpin results need to be correlated – and then data from the device driver can be delivered. All of this activity represents tens of thousands of commands and calls – and results in a lot of processor management/communications overhead just to manage device interaction with a processor.

In 2008, Intel introduced a new bus known as QPI I/O model that offers a direct mechanism for I/O devices to communicate with processors – without all of the overhead associated with traditional I/O driver calls and device management. Intel's QPI is a point-to-point processor interconnect that allows components to connect to I/O hubs or routing hubs on a network. And, in 2014, IBM introduced a new version of its POWER processor known as POWER8 – and this processor supports a new channel know as Coherent Accelerator Processor Interface (CAPI) – which is essentially a hardware-based I/O accelerator.

Both QPI and CAPI make it possible to attach terabytes of Flash memory to their associated processors. Each interface is faster than using a traditional I/O channel (because of the reduced overhead described above) – and, as a result, Flash memory can be deployed as a separate "memory layer". The reason an enterprise might want to exploit one or both of these interfaces is that Flash Cache costs significantly less than main memory (in one case, a system using more expensive memory DIMMs costs 3.2x more than a similarly configured system that uses Flash memory to run the same workload). Readers should note that DIMMs are faster than Flash memory – but trying to put all data in main memory is prohibitively expensive for most enterprises. Using Flash memory, enterprises can now afford to build less expensive but equally functional in-memory systems environments.

> ***EMC is positioned well in both markets. Its XtremIO Flash arrays have already claimed the number one spot in the enterprise Flash array marketplace. And its DSSD acquisition has positioned EMC as a Flash-as-memory vendor. We can clearly see opportunities for EMC to become a major player in this market as a supplier of large Flash memory devices/configurations to major systems makers as well as systems integrators.***

A Closer Look at EMC Flash-based Storage Offerings

EMC's Flash-based storage offerings include:

1. All-Flash arrays including XtremIO and VNX-F (a dense all-Flash array); and,
2. Hybrid HDD/SSD arrays including VNX, VNXe,VMAX, and Isilon.

Further, EMC has been fairly aggressive in software defined storage (as we describe in this report) – making it possible for buyers to add even more functionality and to improve the performance of its storage offerings.

XtremIO is EMC's flagship all-Flash array – offering high performance Flash storage, and its own operating environment and management tools. The array itself is a scale-out cluster design that allows capacity and performance to be easily sized– and that scales linearly (so adding more XtremIO "X-Bricks" does not degrade performance). This design features no single point of failure

and in many cases tolerates multiple points of failures. The XIOS operating environment has been designed to make sure that XtremIO clusters automatically (without operator administration) perform in a balanced fashion. The XtremIO management environment speeds configuration and deployment – and helps administrators automatically provision storage quickly. Further monitoring tools with real-time metrics – as well as reporting tools that show the savings that are being achieved through thin provisioning and deduplication – are also available.

EMC's VNX-F (the "F" is for Flash) environment provides a low cost-of-entry into the all-Flash market. It is intended to be used for targeted workloads that don't benefit from advanced data services.

As for hybrid solutions, EMC offers its VNX, VNXe, VMAX, and Isilon solutions. All VNX arrays can be flexibly configured with all-Flash or as a hybrid flash (SSD/HDD) drives. Likewise, EMC's VMAX line (10K, 20K, 40K, 100K, 200K, and 400K) can also be configured as hybrid flash solutions. Further, EMC offers "Flash PowerPacks (pluggable SSDs) that make it possible for EMC customers to easily and economically add Flash to their existing arrays.

As for software defined storage, EMC provides "ScaleIO" that helps facilitate convergence, scale and elasticity – while also accelerating SSD performance. And the company's XtremCache can be used to accelerate cache performance for certain applications.

EMC's Flash-as-Memory Offerings

EMC expects to begin shipping Flash-as-Memory offerings next year (2015) using technology acquired through the acquisition of DSSD, Inc. Note that EMC led the Series A investment in DSSD, Inc. – and had been a very active partner in the development of the DSSD product design.

The purposed of this acquisition was to get EMC into the Flash-as-Memory market as quickly as possible. DSSD offered a rack-scale flash storage design that is well suited to serving the needs of applications that need more memory (for example, as a Flash cache, or as a complete Flash store). As EMC described its intentions when acquiring DSSD, the company claimed that "the prospects of what EMC and DSSD can achieve together are truly remarkable. We ventured out to create a new storage tier for transactional and Big Data applications that have the highest performance I/O requirements. Working together with EMC, DSSD will deliver a new type of storage system with game-changing latency, IOPS and bandwidth characteristics while offering the operational efficiency of shared storage. "

The way we see it, this new generation of Flash-as-memory offerings will help drive the deployment of a new generation of affordable in-memory database servers (as we state in this report). The software and databased used in these IMDMS (In Memory Database Management Systems) environments will come from the traditional large database makers such as Oracle, SAP and IBM – as well as from relatively new IMDMS providers and NewSQL/NoSQL providers, such as including Altibase, Aerospike, Exasol, GemFire, Kognito, McObject, MongoDB, NuoDB, ParStream, Quartet FS, VMware and VoltDB. "NewSQL" and "NoSQL" data management market, often differentiated by high-frequency transactions, high-bandwidth, and low latency data access. We also see these new Flash-as-memory offerings opening the door for next generation risk management, fraud detection, high-frequency applications – such as those provided by Pivotal HD and IBM (see this report on next generation risk/fraud management). And finally, we expect EMC's DSSD offerings to serve the needs of enterprise customers who wish to analyze very large Hadoop Big Data databases in real-time.

Customer Comments

EMC provides almost a dozen case studies (on this site) that contain XtremIO deployment scenarios and customer comments. These customers include Baptist Health Systems, Boston Scientific, CMA, E*Trade, Georgia Tech, Keolis, KPIT, and CLAL Insurance – and deployments include consolidation, virtual desktops, Oracle data warehouses, and SAP applications.

Summary Observations

As could be expected of the storage market leader, EMC identified the migration trend from HDDs to SSDs early (in 2005) – and began investing R&D dollars into building all-Flash and hybrid-Flash offerings. It also set aside investment dollars for acquisitions such as the acquisition of XtremIO – and for seeding Flash companies such as DSSD, Inc.

Most noteworthy about EMC's Flash storage strategy is that the company has been sensitive to its customer's needs when it comes to HDD to SSD migration. The company recognizes that Flash can be a disruptive technology – so it has made it possible for its customers to deploy Flash storage in their existing arrays, to purchase hybrid HDD/SSD solutions (an evolutionary approach) – or customers can migrate in-total to entirely all-Flash arrays (a revolutionary approach). Also noteworthy is that EMC has been aggressive in software defined storage – the company is already delivering software defined storage solutions for its products that help speed performance while also offering facilitate HDD/SDD convergence and scale.

EMC's HDD/SSD storage migration strategy is paying big dividends for the company. In less than 1 year, its XtremIO All-Flash storage offering is now running at a $500,000,000 annualized run-rate. Further, The 451 Group and IDC have declared EMC the leader in the enterprise all-Flash storage marketplace for 2014.

As for future Flash growth opportunities, EMC's DSSD acquisition has positioned the company extremely well in what we call the emerging Flash-as-Memory marketplace. This marketplace can be characterized as a systems design market where Flash memory will be used as somewhat slower than main memory DIMMS – but at a much more affordable price point. Flash memory systems designed for this marketplace have the potential to lower the cost of IMDBMS configurations by two-thirds – making it possible for enterprises that previously couldn't afford a huge investment in large main memory configurations to suddenly buy more affordable Flash memory configurations. We believe that there is huge potential for these types of Flash memory based systems – and EMC is very well positioned to become a driving force in this evolving marketplace.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662