



## Outlook Report

# Enterprise Needs for Solid State Drives and Data Centers – A Market Matures

*An Outlook Report from Storage Strategies NOW*

January 14, 2015

By Deni Connor and James E. Bagley  
Client Relations: Phylis Bockelman

Storage Strategies NOW  
8815 Mountain Path Circle  
Austin, Texas 78759  
(512) 345-3850

SSG-NOW.COM

Note: The information and recommendations made by Storage Strategies NOW, Inc. are based upon public information and sources and may also include personal opinions both of Storage Strategies NOW and others, all of which we believe are accurate and reliable. As market conditions change however and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Storage Strategies NOW, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.

This report is purchased by Dell, who understands and agrees that the report is furnished solely for its distribution to customers and prospects, without the prior written consent of Storage Strategies NOW.

Copyright 2015 All rights reserved. Storage Strategies NOW, Inc.



Storage Strategies NOW | [www.ssg-now.com](http://www.ssg-now.com)  
512.345.3850 | [info@ssg-now.com](mailto:info@ssg-now.com)

Sponsor



## User comments on Flash/SSDs

“Flash is the new normal.” – *A major systems vendor*

“Storage is the next big battlefield in technology. Every other major trend in computing today relies upon storage in some way or another. Server virtualization, VDI, Big Data, Cloud, disaster recovery, and even Security Information Event Management, all rely on fast, robust, and easily managed storage. We take flash storage for granted at a mobile level, but as the technology has improved, we are seeing that desktops, laptops and servers are all seeing significant performance and power consumption benefits from SSDs and similar storage. We are finally realizing an improvement in storage performance that matches the majestic gains we've had with CPU performance over the past two decades.” – *CIO*

“Flash has been used for entire environments to meet our business objectives, as well as in tiering to meet general I/O and latency objectives.” – *Storage architect, large healthcare information technology vendor*

“[Flash] really helped solve some historical performance problems and enabled our business to serve its customers in a better way.” – *Server and storage engineering, large insurance and investments company*



## Outlook Report

Dear Reader,

SSG-NOW looks at technologies and products when confusion exists in a specific market. What is solid state storage or Flash Memory? How is it used best? What are the use cases for SSDs/Flash technologies? When should you adopt SSDs/Flash and how? This report will answer those questions about solid state/Flash-based storage and in what situations it is best used. If as a customer, you have further questions about solid state drives or Flash memory, how it is implemented and deployed and where it is best used, please feel free to contact us. We will enjoy talking to you.

Deni Connor

Principal Analyst

dconnor@ssg-now.com

James E. Bagley

Senior Analyst

jbagley@ssg-now.com

- Sponsor ..... 2
- User comments on Flash/SSDs ..... 3
- Executive Summary .....7
- Business Drivers for Flash and SSDs ..... 9
  - Flash management details .....10
  - Why it matters ..... 11
- The State of the Solid State Technology.....12
  - Major market participants .....13
  - Memory Technologies other than Flash.....14
    - DRAM.....14
    - NVRAM .....14
    - MRAM and Spin-Torque.....14
    - Carbon Nanotube.....16
    - Phase Change Memory – Diode Transister Memory..... 17
    - RE-RAM, Memristor ..... 17
  - Conclusions ..... 17
- IT Professionals Adoption Survey of SSDs/Flash .....18
- Flash Deployment Challenges and Solutions ..... 20
  - Over-provisioning..... 20
  - Wear-leveling..... 20
  - Write endurance ..... 21
  - Write amplification ..... 21
  - IOPs – Read/write performance ..... 22
  - Latency ..... 23
- Implementations ..... 24
  - HDD Form Factor Flash SSD Drive ..... 24
  - PCIe Format Flash SSD..... 24
  - All Flash Appliance/Array..... 24
  - Hybrid HDD and Flash Array ..... 25
  - Fabric Connected PCIe Flash..... 25
  - NVMe ..... 25
- Best practices and recommendations .....27
  - Understand the Input-Output bottlenecks .....27
  - Understand the technologies .....27
  - Don't use a consumer-class device in an enterprise application ..... 28
  - Go for the low-hanging fruit ..... 28
  - Evaluate Multiple Offerings and Negotiate features and pricing ..... 28

Enterprise Use Cases.....	30
Overview .....	30
The details.....	30
Why it matters .....	30
Academic Research.....	30
Archive .....	30
Big Data.....	31
Biosciences and Biopharma.....	31
Cloud Compute and Storage .....	31
Finance.....	31
Government .....	32
Internet Content .....	32
Manufacturing .....	32
Media and Entertainment.....	32
Oil and Gas.....	32
Retail .....	33
Surveillance.....	33
Application-specific Use Cases.....	34
Analytics.....	34
Simulation.....	34
Transaction Processing .....	34
VDI .....	34
Server virtualization .....	34
Vendor/Product Profile .....	35
Vendor Name: Dell.....	35
Signs of Maturity .....	37
Our Take .....	38

## Executive Summary

The advent of solid state storage devices, which use NAND Flash technology to provide persistent data retention when systems are not powered up, has made an increasing mark on the enterprise data center. In addition to higher access speeds, NAND Flash provides lower power consumption than its counterpart hard disk drive (HDD) technology.

Over the last four years, enterprise consumption of solid state storage has increased to an estimated \$10 billion in the current calendar year. Storage devices include solid state drives (SSDs), which conform mechanically and with interface electronics identical to hard disk drives (HDDs). These devices can be placed in an existing disk array or server drive bay, and depending on the system, can offer immediate performance improvements. Several implementation strategies are built around PCIe cards that are placed directly in the server and utilize the performance of the server's bus to provide acceleration to applications.

In the last four years, a new class of all-solid-state storage arrays has become available from a number of start-up companies, as well as established storage vendors. There is also a class of storage called a hybrid storage array. These devices use a combination of solid state storage for performance acceleration and hard disk drives (HDDs) to accommodate capacity storage. Hybrid arrays account for some \$14 billion in the current year; this includes both rotating media and Flash.

Technology advancements in NAND Flash have followed Moore's law, decreasing cost per capacity by shrinking the size of the memory cells. The smallest cell sizes currently available are about 19 nanometers. Several array manufacturers have announced products based on this cell size, as well as a smaller 16 x 20 nanometer size currently available for sampling by original design and system manufacturers. Most manufacturing is at 19/20 nanometer cell sizes using multi-level cell (MLC) technology, which stuffs two data bits in each cell. New processes from NAND Flash foundries can build the die in three-dimensional cells to add capacity without further shrinking the cell sizes. Major foundry owners SanDisk/Toshiba and Samsung have announced that they are developing three-dimensional cell technologies, which have the promise of continuing to reduce the cost per capacity of NAND Flash. Since higher die (chip) densities can be achieved in the vertical dimension, further shrinkage in the north-south and east-west dimensions can be limited, which means lower latency and lower wear factors.

The importance of the continued development of NAND Flash capacities per area is inherent to a improving a wide variety of consumer devices, such as smart phones and tablets. As has been the case in most electronics technologies, it is the high volume consumer products that are driving the cost points of products used by enterprises. The use of the consumer products is not free, however, as more exotic hardware and software strategies need to be deployed to

maintain the reliability of the Flash circuits over time.

A new type of Flash deployment within servers has taken hold this past year by the acquisition of Memory Channel Storage technology pioneer Smart Storage by foundry owner SanDisk. IBM, which sold its industry leading market share of major servers to Lenovo early this year was the first major server manufacturer to commit to this architecture and Lenovo is currently shipping these products.

There are new technologies being developed which may someday provide persistent storage at costs that compete with NAND Flash. These technologies include Carbon Nanotube, Magnetic RAM, Memristor and Phase Change Memory. While these technologies will produce persistent memories that perform much faster than NAND Flash, barring an unforeseen breakthrough however, it will be a number of years until these technologies rival Flash at a cost per capacity basis. In the meantime, the manufacturers of these new technologies will be targeting the Non-Volatile DRAM (NVRAM) applications, where data persistence is achieved by adding batteries or capacitors and Flash backup to these circuits, adding additional costs and reliability concerns. A good example here is in the 'Internet of Things' where devices monitoring a wide variety of data points need persistent memory only for a relatively short period of time because of near-real time connectivity. In these applications, Mega-bit storage devices will suffice, as opposed to enterprise storage systems where Giga-bit capacities are required.

The past year has witnessed industry consolidation at an unprecedented pace as Flash foundry owners and major HDD manufacturers have extended their product portfolios in what seems to be a 'keeping up with the Joneses' feeding frenzy. The most recent acquisition of Fusion-io by SanDisk represents perhaps the strongest example of this trend where a large (\$500 million per year) run-rate publicly-traded supplier of PCIe storage and systems was gobbled up to bolster SanDisk's end user and OEM business.

In this report, we focus on the growing reliance of industry on persistent memories and the growing number of deployments, technologies and participants in the field.

## Business Drivers for Flash and SSDs

The primary characteristic of Flash memories is persistence, that is, the ability to retain data through a power outage. While not as fast as the DRAM memories that dominate the CPU memories in all computers, Flash memories are thousands of times faster in random read applications than hard disks. This inherent speed is the primary business driver for usage in the enterprise. While this speed does not come without much higher costs on a capacity basis when compared to hard drives, it often has advantages in the total cost of ownership where transaction time, employee productivity, power consumption and footprint are taken into account. These advantages represent the explosive growth of deployments in the enterprise applications.

Enterprises use solid state storage in both data center deployments as well as client-side storage in laptops, tablets and smart phones. Some data center operators have found out the hard way by deploying SSDs designed for laptop usage in their high-use storage arrays. While things work great for awhile, without enterprise Flash management capabilities, these drives will have continually degraded performance and outright failure due to the basic characteristic of NAND Flash wear. Another contributing factor to both reduction of wear and improved write speed is the use of deduplication and compression algorithms. Because NAND requires an area to be erased before it is written, the reduction of the data being written contributes to storage system performance on a positive basis.

Vendors who have been serving enterprises have developed strategies for the enablement of reliability within solid state storage devices. These include several strategies starting with basic error detection and correction technology, adding redundant capacity, analyzing usage of areas of the Flash memory that may be incurring higher wear and striping data across multiple storage devices and even systems. In addition, the ability to add scale to systems from small capacity to large capacity without replacement or downtime is extremely important in the data center.

Data center Flash users have high expectations for these expensive systems. First, the need for speed in operation is often the primary driver. In cases where the number of transactions that can be processed in a given period of time, the benefit can be easily understood. Other cases, speed is important and related to optimizing productivity of human resources. A simple example is given by a major Internet Service Provider that operates tens of thousands of servers. If a server needs to be rebooted from HDD, the technician overseeing the operation can wait upwards of five minutes to tell if the operation solved a given problem. By simply placing an SSD as a boot drive in each server, the time the technician is waiting is reduced to a matter of seconds. Due to the drop in cost of Flash, these boot drives can be obtained for less than \$80. While a critical enterprise application, a boot drive does not have the same requirements of a drive in a high use storage array. This is an important distinction and again, reinforces the importance of understanding the ultimate use case for the solid state device.

Flash storage can produce high-speed operations at much lower power consumption than HDDs and in a much smaller data center footprint, reducing operating costs accordingly. When engineered and managed properly, Flash storage provides extremely high levels of reliability, especially when compared to failure rates encountered with mechanical media.

## Flash management details

IT professionals expect a three to five year use-life of any system. However, a consumer-grade SSD may be worn out after as few as 500 total drive capacity writes. Five hundred times is plenty for a typical laptop application, but in an enterprise environment ten to twenty drive-writes can occur in a single day. Further complicating the wear issue is that of a particular area on the drive may receive an inordinate number of writes, causing that portion of the drive to become unreliable.

There are numerous wrinkles on the algorithms to distribute wear across the storage device. Commonly used is the virtualization of the physical device so that all writes become sequential, and the entire drive is written before any area is re-written. Strategies of deduplication and compression are addressed in the controller for the drive, or in the case of a hybrid or solid state area, across the entire system. This function is both important to write performance, but not all data reduction algorithms are created equal. The best are those that can handle the function in-line in real-time without affecting overall system performance. That's a tall order. The worst are those that operate after-the-fact.

Those strategies actually impair the system's wear control because the data is essentially written twice, once in bulk and again when reduced. These are important issues to address with any drive or storage system vendor. An important strategy is that of over-provisioning. In over-provisioning, SSD storage capacities are declared on a net basis. A 400GB drive has an actual capacity of 512GB. The difference is used to backfill areas that are wearing rapidly. Each Flash cell may have a different wear characteristic, so the best systems monitor the health at a cell basis.

High availability (HA) is an expectation of the enterprise data center, meaning that no single component failure can cause any system data to become unavailable. This feature requires redundant controllers on a given system, or redundant nodes in a clustered architecture. Other redundant components involve power supplies, fans and even backup power strategies if a large amount of volatile DRAM is involved in the controllers to improve performance. No one ever said enterprise storage systems were easy to pull off. Error detection and correction algorithms are also critical to system reliability. While an individual SSD or Flash module may have a controller doing this on a local basis, hybrid and all-Flash storage systems implement features that extend detection and correction across multiple SSDs, modules, storage nodes and even clustered system components. Each vendor should be able to explain their approach in understandable terms and backup their claims with real-world data. Some systems have hot-swap capabilities and upgrade processes that do not involve downtime. These are difficult to implement and will ultimately contribute to the overall system cost. In all cases, the buyer must be aware of the trade-offs.

## Why it matters

Enterprises have expectations that their data is reliably stored and will be available in the long haul. The most important component of stored data is that which is needed the most frequently and at the highest need, the so-called Tier 0 of storage. Solid state storage is becoming the gold standard for this type of data. Despite the technology issues related to wear-leveling, high availability and data reliability, the enterprise data center can obtain these systems from multiple sources. Many buyers will stick with their long-term storage vendors, at a price, while others will go with younger companies with newer technology and lower price tags.

In any event, enterprises should evaluate companies based on technology, reliability, reputation and staying power. Technology is also an important dimension because often it can be the difference between a cost-justifiable implementation and a system that just is not worth the money and effort.

## The State of the Solid State Technology

**N**AND Flash memories have been in use in mobile devices for about twenty years. The primary advantage of Flash is that the data written to it is persistent. Unlike dynamic random access memories (DRAM), data is retained after the circuit has no power provided. The primary disadvantage is that the memory needs to be written in pages and an entire page needs to be written even if only one bit has changed. Pages vary in size, ranging from 1KB to 32K bytes. Once the entire memory is written, each page needs to be erased before new data can be written. This is referred to as a 'program-erase' (PE) cycle, which takes time. However, that is only the beginning of the problem. Each PE cycle causes wear and tear on the memory and the memory will eventually wear out. For this reason, Flash memories need to have careful care and feeding algorithms. These include processes referred to as 'wear leveling,' which distributes writes across pages so very active areas do not wear out quickly. Most wear-leveling processes start by treating the entire memory as a sequential device, writing pages until the memory has been completely written and then 'rewinding' to start over again. This requires an index of virtual pages to physical pages.

Moore's Law states that we should expect silicon device capacities to double every two years, as the companies that manufacture the devices continue to shrink the internal circuits using ever-smaller unit densities, driven primarily by demand for consumer-class Flash memories. In fact, we have seen this in the last four years, as Flash lithography has dropped from 40 nanometer (nm) units, to 30nm units, to 25nm/24nm, to 20nm/19nm units and this year finally to 16nm/15 nm units. Each reduction allows higher capacity devices at similar production costs. However, smaller cells also require more error correction technology.

A recent round of industry consolidation and partnering has created a situation where the vast majority of Flash memories are manufactured by only three consortiums: SanDisk-Toshiba who operate a gigantic facility in Japan, Samsung who operates similar facilities in Korea (Samsung recently sold its HDD manufacturing operation to Seagate) and Micron. Business is good for these 'foundry' owners. Demand for smart phones and tablet computers have gobbled Flash memories at astonishing rates. These consumer devices create the volume production necessary to support the rapid expansion of enterprise SSDs.

Flash memories come in three basic flavors: single-level cell memories (SLC) that which have a single data bit in each cell, multi-level cell (MLC) that has two bits in each cell and three-bit-per-cell devices that have an inaccurate moniker of 'TLC' (there are actually eight levels in a 'TLC' device). While TLC devices are primarily used in applications, such as memory sticks) that do not require wear leveling, some vendors are looking to these devices to provide capacity cost advantages during the next two years, particularly in client-side applications.

Flash memories provide persistent storage with resilience during the loss of system power and speed of access that provides data in microseconds, a rate that is more than 100 times faster than that available from HDDs. Other advantages include lower overall power consumption and the ability deliver high quality of service levels for a variety of applications. However, Flash memories are more expensive per capacity than the lowest cost HDDs and, while an HDD has virtually

unlimited ability to process data writes, Flash memories are subject to wear that must be handled at both an individual Flash device and on an overall system basis. Flash memories provide rapid performance for data reads but are slower for data writing due to a characteristic that requires the erasure of a physical page of data on the Flash module before it can be written. This slows down the write process unless other system functions can mitigate this step. The performance of all storage systems relies on the use of high-speed volatile memories known as DRAM.

A developing Flash technology involves layering of silicon in the vertical dimension. Often referred to as NAND<sup>3</sup>, or three-dimensional Flash, all major Flash foundries are developing this technology, with Samsung and Toshiba/SanDisk expecting products in the two-year time frame. NAND<sup>3</sup> has many advantages, but also had huge technology challenges to be ironed out.

## Major market participants

Due to the substantial market consolidation in the last twelve months, a different view of the major market participants is developing. More and more, the foundry owners and the historical purveyors of hard disk drives are becoming giant storage vendors providing devices ranging from components to full storage systems.

The following table compares the largest companies based on most recent revenues. Of extreme interest is Micron, historically known for its manufacture of all types of memory devices. Its dramatic growth puts it at the top of this list. Recently, Micron formed a storage division under former Dell Storage general manager Darren Thomas who managed the EqualLogic and Compellent acquisitions while at Dell.

Company	Recent Revenue (\$Billions)
<b>Micron</b>	16,538
<b>EMC</b>	16,132
<b>Western Digital (includes HGST)</b>	15,130
<b>Seagate</b>	13,724
<b>Samsung*</b>	9,832
<b>Toshiba*</b>	6,391
<b>NetApp</b>	6,332
<b>SanDisk</b>	6,170

\* Branded Flash Memory Sales Only

## Memory Technologies other than Flash

### DRAM

Dynamic Random Access Memories (DRAM) are the fastest and easiest to manage. They are the basic computer processor memories used in all devices ranging from a cell phone to a multi-core server. Moore's Law is alive and well with DRAM technology and indeed, these devices continue to get faster and cheaper but do not have the same density per area, as Flash. DRAM, while close to a cost per GB of Flash, needs a larger area on the board and more power. Just like Flash, DRAM memories have the same electron size problem, so a brick wall exists for this critical system component. Since they are dynamic, which means they need to be refreshed frequently, DRAMs need a power source at all times, otherwise they forget their contents.

Despite these issues, DRAM usage continues to expand in high-performance computing environments, with acceleration appliances and even entire databases held within rack-mount devices. Since they have no real management issues like wear-leveling and program-erase latency, applications requiring speed at any cost will use DRAM. Backup batteries are deployed to keep the appliances from losing their contents in applications where data loss due to power failure is not acceptable. New generations of hybrid non-volatile DRAM systems are being deployed using super-capacitors, coupled with a Flash memory circuit to save the contents of the DRAM to Flash in the event of a power failure.

### NVRAM

Non-volatile random access memory (NVRAM) combines DRAM with additional circuitry in order to make DRAM memory persistent through a power failure. Depending on the amount of DRAM that needs to be protected, additional circuitry can involve capacitors, batteries, Flash memory or combinations, such as a capacitor tied to a Flash module. In any event, NVRAM is expensive because circuitry and components need to be added to the already expensive DRAM.

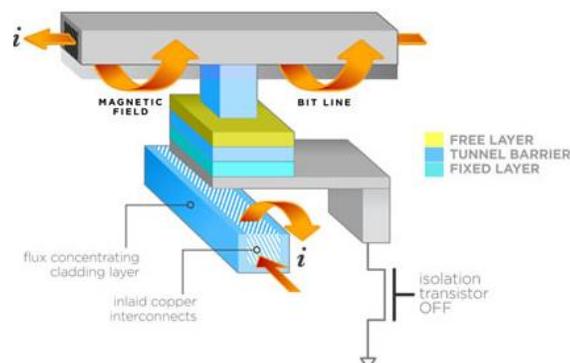
### MRAM and Spin-Torque

Everspin announced that its shipments of magnetic random access memory (MRAM) have increased substantially in the last two years. MRAM combines complementary metal oxide semiconductor (CMOS) technology with magnetic cells inserted between metal layers to create a very fast memory that can retain its state for years without power supplied to it. A simple transistor per cell detects the amount of resistance in order to determine the state of the bit. This technology is referred to as memory tunnel junction (MTJ) because the areas of charge are separated by a tunnel. One wire is used to charge the magnetic cell to a zero or one state, and another wire is used as a constant magnetic field, in a setup referred to as toggle cell design.

## MTJ STORAGE ELEMENT



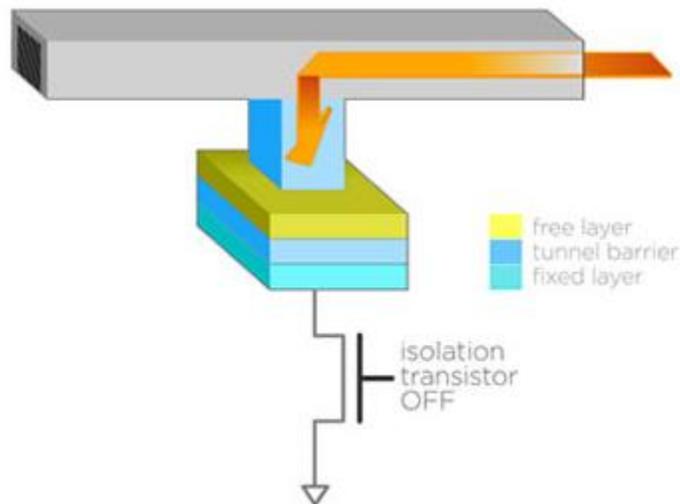
## Toggle Cell



While currently at a fairly large geometry (around 100nm), the advantages of the high-speed memories for processor memories and caching are already creating high-volume demand, which will allow the continued development of the technology. Currently, major customers in the aerospace, energy and automotive fields are utilizing the low latency (<50ns) and wide operating temperature (-40C to +125C) of MRAM in applications that require memories to withstand power failures and, unlike Flash memories, have unlimited read-write cycles and can retain data for more than 25 years.

Everspin has 64Mbit DDR3-compatible components and its product roadmap is clearly headed towards Gigabit densities that will allow MRAM to compete effectively with nonvolatile DRAM in applications that currently require an expensive battery back-up or super-capacitor combined with Flash memory to provide data persistence after a power failure. The key to shrinking the size of the lithography to 40nm and even <30nm is going from the double-wired MTJ 'toggle' set-up to what is called spin torque. Spin torque allows a bit to be programmed directly without the double wire setup. The MTJ still remains paired with a transistor to sense the status of the bit. Since the MTJ requires less current as its lithography shrinks, the transistor can likewise be shrunk. Everspin sees 8Gb and 16Gb parts appearing around 2015. While not intending to price attack the Flash and phase-change memories (PCM) that should be appearing in that time frame, the needs of applications like high-speed non-volatile processor workspace and non-volatile write caches should account for about one third of the total dollars spent for DDR, FLASH and non-volatile solid state memories.

### Spin Torque Cell



Everspin came from Motorola's Freescale business unit in 2008 and is a privately held company. It holds over 600 patents and filings with over 175 already granted covering the MRAM and spin-torque technology. The production steps involve receiving a CMOS wafer from one of the generic foundries and applying the MTJ cells and final packaging at the Everspin Chandler, Ariz. facility. As demand and applications for the technology increase, Everspin expects to license to process to one or more of the major memory fab owners.

### Carbon Nanotube

Nantero is promoting an emerging memory technology built from Carbon Nanotubes (CNT) that it calls NRAM. Nantero has been developing the technology since 2001 and has a patent portfolio of 165 issued and another 200 patents pending. Rather than developing foundries to produce the components, Nantero has focused on licensing the intellectual property and providing technology support to foundry owners and fab-less silicon producers. Complementary metal oxide semiconductor (CMOS) fabs are currently capable of producing the parts with recipes provided by Nantero and its supply chain partners.

CNT works like a mechanical switch -- when the carbon tubes touch, the switch (bit) is set and when they are separated the switch is reset. This provides an extremely robust data persistence that has been tested for retention for decades at temperatures in excess of 300°C. Further good news is the very low current necessary to set and reset the circuit. The devices under development by Nantero's licensee's have already produced 4Mb sample parts that operate faster than DRAM, have no endurance issues like flash and operate at 100 times the speed of flash. The processes to bake the parts in the foundry are straightforward and materials used in the nanotubes are available from the typical fab supply chain. Not, that this has always been the case. The initial materials used to create nanotubes had too high a concentration

of contaminating metals, which can destroy a CMOS fab. Several years of incremental refinement have produced the competing technologies of magnetic RAMs (MRAMs) and phase change memories (PCMs).

## Phase Change Memory – Diode Transistor Memory

A new type of device that solves the electron size problem in DRAM and Flash memories is becoming a reality. Phase Change Memory (PCM) uses heat to change the conductance of an area on the circuit. Since conductance, rather than electronic charge, is used to determine the state of a bit, the lithography is not limited by electron size. This allows a memory that can be directly written without a page PE cycle. These heat cycles are very quick and the wear caused by a cycle is small, each bit can be written a billion times. While small, the wear is enough, in terms of enterprise computing, to require a wear-leveling process. The technology belongs to foundry owner Micron as the result of Micron buying Intel out of a joint venture called Numonyx. The status of Moore's Law, in terms of capacity and cost per bit, is lagging Flash by years, as Flash continues to be developed with solutions that will keep it viable for many years.

Contour Semiconductor, a well-financed startup, is currently developing stacked diodes using PCM architecture.

## RE-RAM, Memristor

Resistive RAM is a technology under development by startup Crossbar Technology, which has the possibility of scaling very large-scale devices.

Memristor is a promising technology that can use resistance levels as the representation of bits. The technology is still in the research labs of major organizations including HP, IBM and Samsung. The devices remain to be commercialized.

## Conclusions

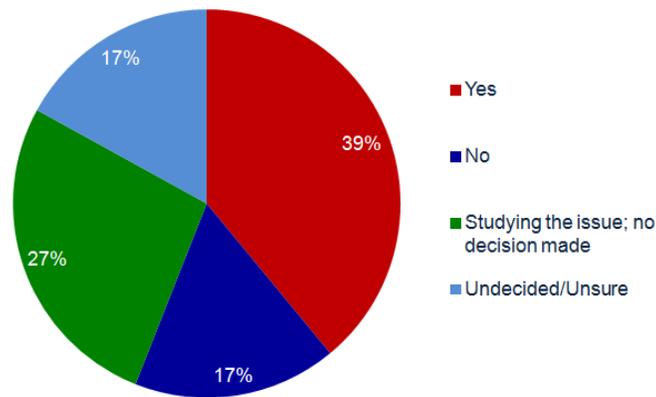
While there is considerable technology under development, barring an unexpected technology breakthrough, there will be no alternative to Flash memories in the five to ten years, perhaps longer. The reason is simply that the Flash foundry owners continue to invest much more money in the development of new Flash production technologies than what is available to developing alternative technologies.

## IT Professionals Adoption Survey of SSDs/Flash

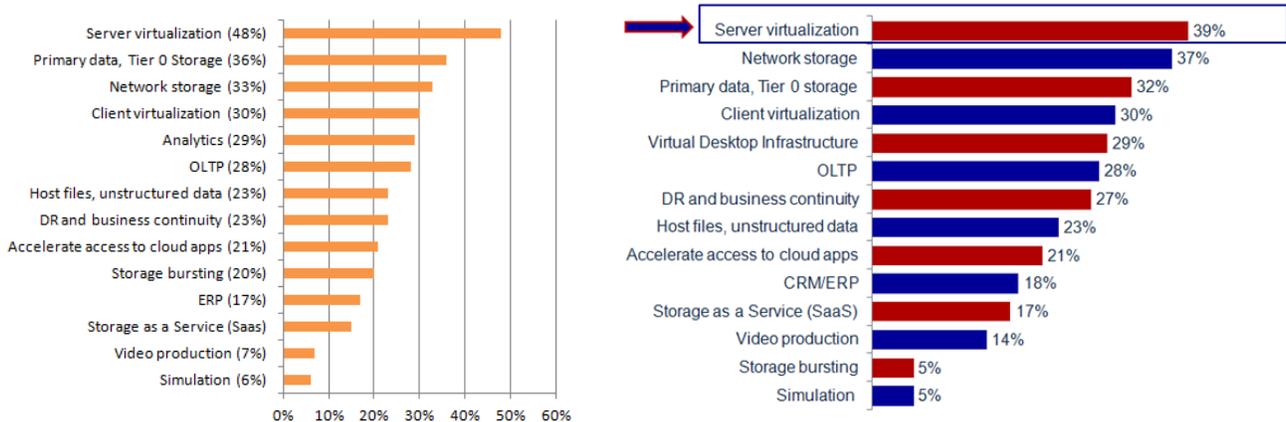
Storage Strategies NOW and ITIC conducted a survey of IT professionals on SSD/Flash adoption and deployment in 2013 and 2014. Among more than 550 respondents in 2013 and 587 respondents in 2014, the results were balanced across the enterprise. In 2014, 36% of respondents were from SMBs, 28% from mid-size and SME environments and 36% represented large enterprises. As many as 47 vertical market segments were surveyed.

Of the results, 39% of the respondents currently use or have definitive plans to deploy SSD/Flash technologies.

Does your firm currently use or have definitive plans to deploy Flash/SSDs?

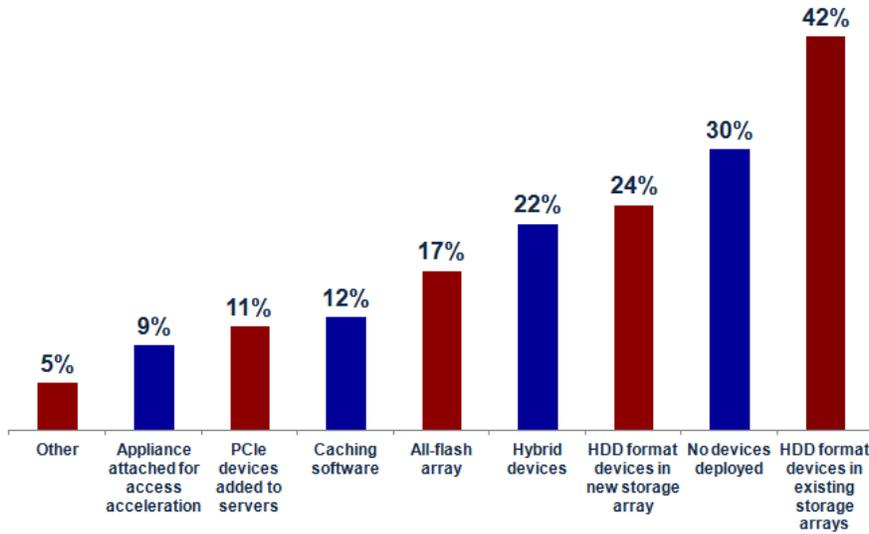


In 2013, 48% of the respondents said they would use SSDs/Flash for server virtualization; 30% for desktop virtualization; 29% for analytics and 28% for OLTP. The responses remained much the same in 2014.



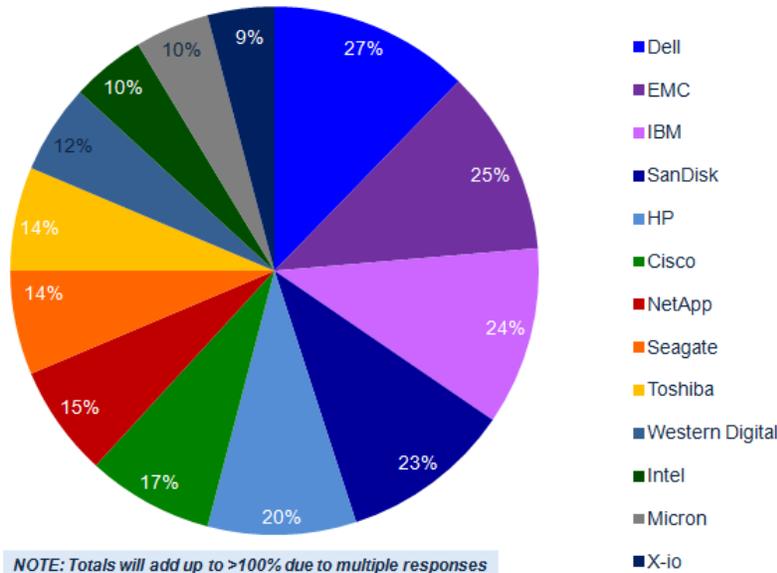
The survey also showed that by a large percentage (42%) will use SSDs/Flash instantiated in existing storage devices. The remainder will purchase new SSD-based storage (63%), add PCIe adapters to servers (11%) or use appliances to accelerate performance (9%). Twelve percent (12%) will use caching software to accelerate performance.

**What type of solid state drives or high speed memory storage devices has your firm deployed or definitively plan to deploy? (Select ALL that apply)**



Respondents identified Dell as the number # 1 leader in Flash/SSD adoption, leading over EMC and IBM. Companies such as PureStorage and Nimble Storage each accrued a 4% market share.

**If your firm currently uses/definitively plans to deploy Flash/SSDs what vendors do you use or expect to use? (Select ALL that apply)**



## Flash Deployment Challenges and Solutions

In spite of SSDs/Flash becoming a mature, well-adopted market for tier 0 and tier 1 applications, concerns still hover of certain features of the technology. Among these concerns are over-provisioning, wear-leveling, write amplification and write endurance.

### Over-provisioning

Ever since the late 1980's, HDD capacity has been stated in terms of decimal GBs, rather than the higher binary equivalent. That's because a KB is really 1,024 bytes. A binary GB is  $1024^3$  (1,073,741,824). A decimal GB is 1,000,000,000 bytes, a difference of nearly 74MB, or 7.4%. Where do the lost bytes go? They are used for error correction, housekeeping data and bad blocks that are the result of tolerances in the manufacturing process. They are used internally by the drive controller, but are not available for storage of data or metadata. With Flash memories and the wear factor of multiple writes and, to a lesser degree, multiple reads, even more of the raw capacity of the memories are reserved for controller use. These uses include error correction coding of 128 bits, 256 bits or even more for every page, which is typically around 4 kilobytes. Some file systems add information to the blocks and a problem related to the alignment of blocks on the physical Flash pages can cause wasted space. In addition, as blocks become weaker they get put in a garbage can. This can be detected by a use map coupled with increased error correction activity. RAID-like striping of data across multiple Flash dies (chips) increases reliability and error recovery, but adds cost in useable space.

Consider the current crop of SSDs that come in 100GB, 200GB and 400GB capacities. The raw capacities of these devices, in binary, are 137,438,953,472 bytes, 274,877,906,944 bytes and 549,755,813,888 bytes, respectively. Instead of just a 7.4% per GB over-provisioning that has become the norm in the HDD business, over-provisioning in the SSD business can be as high as 37% per GB or more. We explore the use of this space below in the discussion of write endurance and wear leveling.

Over-provisioning is a necessary evil with Flash devices. The amount of extra space has a negative effect on the cost per GB, but a positive effect on the long-term use life of the product. To some controller strategies, the knowledge of a large expanse of extra space makes controller development easier. In others, particularly as competition increases, the luxury of oceans of extra pages will demand improved controller strategies and improved Flash memory components. The reporting of raw capacity versus useable capacity has improved in recent years, but it remains an important question to ask the purveyor of any Flash-based product. Again, over-provisioning doesn't apply to DRAMs where the raw capacity is equal to the useable capacity.

### Wear-leveling

The first thing we learned in the mobile device field when we started using Flash was that software could break hardware. If you continue to erase and re-program the same page, the page will wear out. I remember a near painful product recall that dealt with this feature of Flash memory. Fortunately, we found the problem during beta testing and were able to fix it with software upgrades. It did delay the product launch, however, because we had to implement a more intelligent file system on the product.

In the modern SSD world, wear-leveling has become a science and all vendors continue to improve their algorithms in close formation with strategies for correction of errors and studying of use maps. The obvious first step is to create a physical to virtual map of the drive in order to ensure that regardless of what the operating or file system is trying to do with the drive, all pages get used once before they get used again. Further discussion of this is covered on the subject of write amplification.

In addition to the physical to virtual mapping, another type of mapping is used. The number of times a page has been written is usually maintained and due to disturbance of adjacent data, the number of reads is also maintained as a way to determine the health of a page. Since some circuits are stronger than others, a pure count is usually not enough to assure the health of a page. Hence advanced forward error correction techniques, that can be used to correct an error after detection and background processing called 'patrol reads,' are all deployed as part of a wear-leveling algorithm.

### Write endurance

Write endurance in Flash memories, that is, how long a page can remember every bit, will vary greatly with the technology (SLC vs. MLC) and the number of times the page has been erased and re-programmed and even the number of times the page has been read. This is because, over time, the electronic charge of a given bit can degrade to the point of unreadability. We have discussed the strategies to minimize the effects of degrading write endurance. First, a wear-leveling strategy must be deployed that evenly spreads the write activity across the entire number of pages on the device. Second, an error correction coding, involving advanced error correction and a large size, or strength of correcting codes included in each block of data. A RAID-like algorithm that stripes data across multiple Flash circuits is deployed in many devices. While this increases the over-provisioning, it also makes recovery from serious errors possible.

### Write amplification

Write amplification is a 'feature' of Flash memories that impedes random write performance. The basic problem is this: whether you are updating a single bit or more, an entire multi-page block needs to be erased and completely re-written. This takes time, or latency, which can negate the very reason the for the price premium paid for an SSD. Blocks in high-end Flash chips are 256 KB or 512 KB. Another feature of this issue is that brand new drives will operate much faster until the entire space of the drive is written. Then the problem rears its head, unless some serious rocket science is applied. Since erase-program cycles wear out the chip over time, this is a critical issue when developing enterprise controller technology. There are multiple strategies to deal with this and the very bright engineers and scientists working in the field are continually improving these processes and running to the patent office to protect their intellectual property.

The very first thing we learned in the mobile field, where Flash memories have been deployed for nearly two decades, is that you treat the entire memory like a sequential tape. You write every page before you 'rewind' to the front of the tape and start erasing and programming. This is really a case of 'the more things change the more they stay the same.' Many of the algorithms for sequential writing of SSDs originated in such obscure devices as the Digital Equipment 'DecTape' drives, which could be randomly read and written. They were fun to watch (but not so fun to wait for) while they moved forward and backwards to access a file. A similar issue occurs with a highly fragmented hard drive, hence, the venerable Windows

de-fragmentation utility. Whether you are writing a bit or many bytes, the next free page is written and indices are updated to reflect the real data location.

Clearly, if writing Flash is a bottleneck, a good strategy is to reduce the amount of data to be written. Here we have significant compression and de-duplication technology to utilize, even at the chip level, to reduce the amount of data that is written. As any compression and de-dupe guru will tell you, the in-line, primary storage compression algorithms are limited to the data stream being received, unlike the original backup and recovery applications, there is still significant reduction available, but the results can vary widely based on the type of input data. For example, MPEG and JPEG files are already heavily compressed and therefore, trying to reduce them will just burn CPU cycles. The only real way to reduce them is to uncompress them and recompress them using a better algorithm. By the time you do that, you would have been better off just writing the pages. However, controllers that do not attempt to compress and de-dupe in-line will suffer performance issues for typical files like PowerPoint's, Word, database and executables.

Another important strategy involves background processing. Once the disk is full, known obsolete pages can be erased and ready for programming again by a background process. There are other advantages to these background processes in managing Flash that include 'patrol reads,' which can determine if the wear issues are impacting certain pages and other housekeeping important to the health of an SSD.

Benchmark developers have addressed write amplification in performance evaluation of various manufacturers' SSDs by recommending that first the entire drive needs to be written. Other strategies involve writing a mix of compressible and incompressible data. One freeware benchmark only writes incompressible data, causing certain manufacturers to add disclaimers to their data sheets and providing provisos in their datasheets for compressible and incompressible files.

Whether the Flash memory is in an HDD form factor, or on a PCIe card, or in a purpose-built appliance, write amplification management needs to be addressed in the software driver, controller or a combination of both. How write amplification is addressed has significant performance and use life ramifications and needs to be evaluated for enterprise applications.

## **IOPs – Read/write performance**

The primary reason to use solid state memories is to increase the performance of a given computer function. Measurements of performance are touted as features of a given product. There are two important metrics used when measuring performance: Input-Outputs per second (IOPs) and bandwidth, expressed in bytes per second (Bps). Since read performance and write performance can vary greatly for reasons of write amplification and program-erase cycles, IOPs must be taken in context with type of operation, be it read or write or a combined workload and the size of the data being exchanged. We try to get the manufacturers to indicate this whenever stating IOPs in a specification. Typically, IOPs are specified for random reads of 4KB and random writes of 4KB. Bandwidth is usually specified as sequential reads MB/sec and sequential writes MB/sec.

## Latency

Latency refers to the total amount of time it takes to complete an operation, so just because you have a fast device, it does not mean that you can take full advantage of its speed. This is especially true of SSDs. Generally speaking, the farther from the CPU accessing the device, the more latency exists. Thus, DRAM that makes up the CPU's main memory has less latency than sitting on a PCIe card, which has less latency than an SSD sitting on a shared storage network. HDDs have lots of latency doing random reads, because of the time it takes to move the read head to a new spot. But HDDs have less latency during a long sequential write as long as the data being written can keep up with the rotation of the drive. While SSDs will generally improve performance, how much improvement often depends on the application's use of the data and other overheads associated with its access.

## Implementations

### HDD Form Factor Flash SSD Drive

HDD form-factor SSDs represent the highest volume devices in the market because they are deployed in both enterprise data center applications as well as in client devices. While an SSD may be installed in a server, it does not necessarily need to be of an enterprise-class drive. A good example is the boot drive.

Far beyond boot drives, the major deployments for enterprise products are in SANs, network attached storage (NAS) and direct attached storage (DAS) systems. As a bet, we vote for foundry owners taking more and more of the market through both direct technology development and acquisitions.

The drive format products have a great advantage of being able to plug and play into an existing array. This does come at a cost over other implementations. Drive format products need to pretend that they are really a disk drive. This means that an industry standard interface, be it Fibre Channel, Serial ATA or Serial Attached SCSI needs to be built into each drive. Besides being able to plug-and-play into a number of existing products (drive bays, JBODs, storage area networks), the drive format products provide reliability benefits in terms of no single point of failure and compatibility with in-use protection and self-healing strategies.

Foundry owners continue to use their unit cost advantages to compete with other market entrants. There has been considerable consolidation in the HDD form factor market. HDD manufacturing giant Western Digital, while not a foundry owner, certainly has the purchasing clout to keep the foundries competing with each other. WD has purchased SSD maker sTec and caching software maker VeloBit this year. Prior SSD investment by WD include its acquisition of Hitachi GST. Foundry owner SanDisk has added SSD maker Smart Storage to its portfolio this year along with making strategic investments in several startups.

### PCIe Format Flash SSD

While taking advantage of the full bandwidth of the CPU bus, PCIe format SSD products have some interesting trade-offs. Once the exclusive domain of Fusion-io, PCIe products are supplied by a large number of vendors, including foundry owners Micron, Samsung, SanDisk and Toshiba. These devices provide extreme performance at premium price points. They certainly keep to the rule of taking the highest speed devices closest to the server CPU. In many situations this direct-attach deployment makes perfect sense. However, as in all things storage, there are tradeoffs. A big one, especially when you are dedicating a single piece of expensive hardware (these PCIe devices can run to tens of thousands of dollars) to another expensive device (a high-end server), lots of concerns about the expense of redundancy and the potential for single points of failure taking very costly components offline are associated with these devices.

### All Flash Appliance/Array

Storage appliances, sometimes called storage arrays, are purpose-built devices, usually in rack mount form factors, which act as either a shared block device or a network attached storage file system. These devices can be built-up from available HDD format SSDs, DRAM memories, custom Flash modules or combinations.

Appliance interface options typically include Fibre Channel, InfiniBand, and Ethernet physical connections. For storage area network (SAN) block-oriented connections, Fibre Channel Protocol over Fibre Channel, SCSI Remote Protocol (SRP) over InfiniBand, iSCSI over Ethernet and recently ATA over Ethernet (AoE) are popular options. Network attached storage (NAS) options typically include Network File System (NFS), Common Internet File System (CIFS), or Server Message Block (SMB) protocols running over Ethernet. Devices range in size and capacity from a single rack unit (1.75" tall, 1U) to an entire rack supporting hundreds of TBs of solid state memory. Advantages of these devices include their shareable nature, wide bandwidth and scalability. Disadvantages include high centralized costs and single points of failure for many designs. The best uses generally involve applications where an entire dataset can be contained in an appliance or appliance pair and shared across a fabric to a number of servers.

## Hybrid HDD and Flash Array

Hybrid arrays incorporate both Flash devices and conventional HDDs. The combination allows a lower cost per capacity by using the HDD for data that is accessed less frequently and using the Flash components for hotter data. Hybrid arrays incorporate advanced software algorithms for tier management, caching, data reduction and data reliability by using the higher capacity HDDs for mirroring and data striping to avoid data loss due to a single component failure.

## Fabric Connected PCIe Flash

Using PCIe-based Flash memories to accelerate server operations has become a regular practice within data centers of all sizes. Initially the use was no different from direct attached storage (DAS), just a lot faster than using a hard disk drive (HDD). This worked as long as the capacity of the data set could be kept on the Flash card. But to handle larger data sets, typically due to virtualized environments, larger capacities of Flash cards were needed at increasing cost, creating an expensive single source of failure.

The use of write-through caching algorithms became a partial solution to the capacity issue. By keeping hot data on the Flash card and going to shared storage for everything else has become a standard practice. We call this server-based caching version 1.0. While this implementation works well within a single physical server, when a virtual workload is moved to another physical server, the cache is unavailable to the moved workload and the hot data must be re-created, causing a period of impaired performance until the hot data is created on the second physical machine's cache. This behavior can drive an administrator, trying to load-balance a cluster of servers, to the brink of insanity. Another challenge with 1.0 caching is the ability to support clustered applications. Lastly, administrators need to mess around with configurations and scripts to make sure that data among multiple physical servers does not get corrupted accidentally.

## NVMe

Non-volatile memory express (NVMe) is a protocol used over PCIe that is powering a new generation SSDs that can be in multiple form factors, including stand 2.5" HDD form factor devices, but can also be much smaller. An important

feature is that the protocol supports hot-swappable devices, which heretofore has not been available in the PCIe card format devices which require the server to be dismantled in order to swap out a card.

NVMe products are currently shipping from a number of vendors with mass availability in the first quarter of 2015. In addition to the hot-swap characteristic, these devices offer a new level of standardization in the class and higher speeds than the 12Gb Serial Attached SCSI standard.

## Best practices and recommendations

### Understand the Input-Output bottlenecks

While it seems intuitive that slapping an expensive solid state memory device in place of rotating media will always improve performance, the amount of performance improvement may not be dramatic or justify the costs. Conversely, solid state memory devices should not be measured against rotating media on strictly a cost per GB basis. The cost per IO is more important to many applications, which are not storage bound, but are performance bound.

Running virtualized workloads is often the culprit in creating IO bottlenecks. Even workloads that have been optimized for sequential IO performance get their IO requests all scrambled up with other virtual machines running on a server. This is called the 'IO Blender' effect.

Increased access time, or operation latency, occurs the farther the memory is away from the CPU core. Devices attached to the CPU bus have less latency than devices accessed over networks. This is why several vendors we have discussed in this report and others still operating in stealth, are working on what we refer to as 'server-centric' acceleration, where PCIe or direct attached SSD is used as a cache to accelerate access to other file or block storage systems. These solutions can often off-load an overcommitted storage system and avoid a costly replacement or upgrade cycle.

### Understand the technologies

You don't need to be a storage administrator or rocket scientist to understand the basics behind the access protocols, device implementations and solid state device technology. Like all things in our industry, there are multiple dimensions to price, performance, use life and reliability that must be viewed in context with the application. Make the vendor explain the price performance advantage of their technical approach. If they can't that is a huge red flag.

Data reduction techniques of compression and de-duplication provide advantages in capacity used, improving speeds of write processes and limiting wear and tear on the flash. These techniques do have an impact on the storage control speed, however. Be certain to understand the effects on system performance in addition to write speed and capacity used. While much of this report has focused on NAND flash devices, DRAM remains the highest performing technology that does not wear out. But it is not persistent, meaning that when power goes off the device, the data goes with it. This means battery backup systems may need to be deployed, *which wear out over time*. For many users, SLC NAND flash is the best deal, in terms of performance and longevity, but costs are double or more per GB over MLC flash. Fibre Channel (FC) and Serial Attached SCSI (SAS) devices have a faster full-duplex interface than Serial ATA (SATA) devices, which typically cost less. Finally, caching and tiering, while providing improvements in performance in hybrid (both SSD and HDD) applications, will never be as fast as all-solid state devices.

## Don't use a consumer-class device in an enterprise application

Stories of arrays crashing because the SSDs have worn out at the wrong time abound. The important words here are caveat emptor. Know the technology and warranty of the devices being deployed.

Premium SSD pricing is often used by major systems vendors who charge a premium for SSDs but offer 'one throat to choke' and free replacement of components under warranty or support policies. Again, caveat emptor. Know what the premium is. This usually can be checked by interrogating online SSD vendors like NEWEGG.COM.

All enterprise SSD vendors represent warranties in terms of years and 'full drive writes per day,' meaning the capacity of the drive can be written x many times per day for the warranty period.

## Go for the low-hanging fruit

Organizations need to evaluate the time that their personnel spend waiting for operations to complete. The productivity of a workforce can be improved by adding SSDs as boot drives in servers and in storage drives in user workstations and laptops for knowledge workers.

AOL uses SSDs as boot drives in all of its some 50,000 servers. This really makes sense. Keep in mind that boot drives can be set up as nearly read-only devices. That limits the wear and tear factor dramatically. With HDD boot times on the order of magnitude of five minutes and SSD boot times of about ten seconds, the payback is extremely quick. High priced network engineers don't need to sit for five minutes before they find out if a reboot has fixed a server or VM crash. Foundry owner and SSD manufacturer Intel has deployed SSDs in all of its employee laptops, practicing what it preaches and justifying the capital expenditure on improved employee productivity.

If a server is having problems accessing shared storage, the use of new, server-centric caching products can provide an instant payback and reduce stress on the shared storage.

Server-side caching can greatly increase the number of VM's that can operate on a given physical machine, which reduces license fees and operating expenses.

## Evaluate Multiple Offerings and Negotiate features and pricing

There are number of storage system offerings that have Flash memory deployed to increase speed of operation. If possible, get the vendors to compete on an apples-to-apples basis. An understanding of the technologies and underlying costs, as outlined in the beginning of this report, will avoid being 'blinded by science.'

By understanding the effect of Moore's Law in the next few generations and planning storage upgrades to synchronize with the availability of the next wave of lithography size can save enterprises millions of dollars in capital expenditure. While Flash cannot compete on a cost per capacity basis with rotating memory, the operational costs are not represented by this metric.

SSDs don't have to spin up and spin down to save power. They operate at a much lower power per IO than rotating media. A quick look at the first wave of Storage Performance Council energy benchmarks shows that SSD enabled products dramatically lower the annual electric bill in a data center or server closet.

Whenever discussing SSDs we focus on the wear-and-tear associated with program-erase cycles. But when properly engineered into a fault tolerant control process, SSDs are incredibly reliable when compared to rotating media. That factor is shown in the early adoption of SSDs in military and rugged applications. SSDs are not affected by drop shock that would destroy a rotating device and harmonic minimization (necessary in large HDD arrays) is not a factor.

## Enterprise Use Cases

### Overview

The use of solid state storage can be generalized by benefit. The primary benefit of speed can be applied in many use cases. The most common involve virtualization of servers and workstations. Speeding up a complex operation, such as database analysis is another speed benefit. Operational costs are reduced because solid state devices use less energy per transaction. Productivity of personnel is improved because functions complete faster.

### The details

While generic to virtualization, online transaction processing and high performance computing, in this section we analyze use cases within vertical applications spaces where solid state storage is used in ways to increase productivity and speed processes.

### Why it matters

Persistent high speed storage is being deployed by organizations of all sizes because time is money in nearly every human endeavor. Reduction of process time, increased employee productivity and reduced operating costs for a given function are the benefits accrued to many use cases.

### Academic Research

Academia is deploying virtualization at both the server and workstation level. One University found that the reduced IT support time was so beneficial that it could provide the virtual desktop infrastructure (VDI) to replace a conventional PC workstation and save the using department hundreds in operational usage fees. The trick with VDI is to keep the shared storage costs low while providing a user experience that is the same or better than desk top computing. Flash arrays provide the shared storage speed to support implementation. Flash storage is utilized throughout the high performance computer systems utilized in academic institutions.

Reduced run-time in analytical operations and database management are use cases that benefit the Academic environment. Processor and data heavy operations are improved by the speed of indices and meta data access available from solid state storage.

### Archive

Archived information does not have the rapid access demands that top tier data requires, but the speed of Flash storage is highly beneficial to the processing of data for archive. Systems that compress and de-duplicate archive data use the speed of Flash storage to reduce data ingestion time.

## Big Data

The importance of high speed storage to Big Data is the ingestion processing and metadata creation and access. Object Stores have the ability to keep metadata separate from data stores and can place metadata that has high future access priority in solid state media.

## Biosciences and Biopharma

The Biosciences and Biopharma applications require large data sets that require analysis to determine if an expected outcome occurred and to retain these large datasets over time on an ever-increasing basis. While it is not currently practical to maintain these data sets on solid state devices, the use of server-side caching algorithms can greatly speed the completion time of these processes. Similar to Big Data analytics, the metadata involving these large file and database elements can greatly speed the analytic processes when stored on high-speed memory.

## Cloud Compute and Storage

The prime drivers that create enterprise demand for cloud services are driving increased usage of Flash storage in the major cloud service providers. More and more service providers are finding differentiation and revenue opportunities by allowing users to provision workloads with a given quality of service in terms of input-output operations per second (IOPs). If a using organization has workloads that require a given service level, they are going to gravitate to the service providers that can allow provisioning of systems that will support the enhanced performance of SSDs. Major service providers in the Infrastructure-as-a-Service (IaaS) field are allowing customers to provision bare-metal configurations that include server-side SSDs as well as shared storage that includes a varying amount of SSDs. On-ramp providers utilize solid state storage to provide a buffer between a customers' data center and the remote cloud site.

## Finance

Financial services are the epitome of the time-is-money adage. Transactions per second (TPS) is the metric that ties directly to the financial performance of an institution in many applications. Server-side cache and all solid-state arrays are deployed in these applications that can change the economic picture of an organization, and financial services firms have been the early adopters of Flash storage for more than a decade.

Large-scale analytic simulations are the domain of these firms and the speed of outcome of a given simulation is an important component to the feedback loop that dominates organizations of all sizes.

## Government

Military and Intelligence operations require high speed analysis of vast quantities of information. Ingestion of three-dimensional graphical data from a wide variety of sources is a critical use case where solid state storage is used extensively.

## Internet Content

Streaming applications are perfect for Flash because they are read intensive. Examples are video-on-demand and edge-servers delivering promotional content. Web-scale organizations use Flash in servers and direct attached storage for rapid boot and caching, as well as application acceleration and higher virtual machine to physical machine ratios.

## Manufacturing

Enterprise resource planning (ERP) is the critical application for manufacturers and whether operating SAP, Oracle or another software suite ERP performance is greatly enhanced by Flash storage at the server and shared storage level. In fact, major ERP supplier Oracle provides storage systems that are performance tuned to the application and utilize Flash storage in the server, storage controller and storage array.

Workstation virtualization is deployed in manufacturing application because many of these systems are task-worker oriented and have a very thin workspace profile, particularly suited to solid state storage.

Online transaction processing, particularly order entry, is deployed in most ERP systems and customers become increasingly impatient if the order entry process is slow. Workstations used for Computer Aided Design are excellent targets for SSDs because of the productivity improvement allowed for expensive design engineering resources.

Manufacturing data needs to be retained for years and retrieved quickly in the case of a recall or quality problem. This is an ideal use case for tiered storage with the most recent data retained in solid state storage and older data descending in priority to lower cost storage.

## Media and Entertainment

Production and post-production applications are enhanced by the use of solid state storage because they involve expensive labor resources applied to very large data sets. Most labor-intensive operations in production are short-lived but the data is retained for the long term for re-purposing.

Streaming data, similar to web content, is a perfect application because it involves multiple random reads after a single sequential write, very well suited to Flash storage.

## Oil and Gas

Analytics that operate against of four-dimensional database (time included) is standard procedure for determining the location of oil and gas resources. These operations are enhanced by server-side cache and solid state tiered data. The

demand for these analytics have been increased substantially by the use of hydro-fracking processes which increase the output of oil and gas from a given formation.

The current boom in US production of oil and gas has been enabled by these analytics, which prior to solid-state storage, often took days to run and now are completed in a few hours.

Other applications include the typical ERP coupled with process control for refiners. As refining resources become more committed due to supply and demand changes, the use of high-speed storage is increasing.

There is a large component of the modern oil and gas industry, which includes distribution management and retail. These applications have the same time-is-money calculation and are enhanced by the use of high speed solid state storage.

## Retail

Major applications in retail that benefit from solid state storage include point-of-sale (POS), inventory management, seasonal workload variance and analytics.

POS systems increasingly are becoming virtualized server-based computing workloads, because the centralized management simplifies many aspects of what are becoming increasingly complex applications. The check-out process in brick-and mortar retail is the single most revenue-throttling factor and is acerbated by fluctuations in seasonal workloads. Inventory management has long been a function of portable computers that use Flash for program and data storage and Wi-Fi and Bluetooth radios for communications.

Seasonal variance, as the result of holidays and changes in inventory for a particular season, create stress for the core processing systems, which have to be built to handle peak loads. Again, Flash storage provides a method to serve these peak loads.

Retailers are increasingly using Big Data analytics to determine seasonal trends and optimal inventory composition. Often the analytics need to deliver the answer within hours if they are to have any business value. Flash storage is increasingly used for storage of the metadata necessary to support this information base.

Web-based retail has all of these characteristics with the addition of online transaction processing stress with customer satisfaction directly related to the speed that orders can be placed, questions serviced and delivery status information available. Tiered storage and server-side caching provide the functionality necessary to provide essential customer satisfaction.

## Surveillance

Ingestion of data from multiple high speed resources is enhanced by processing that uses Flash storage for staging of data.

## Application-specific Use Cases

### Analytics

In the recent survey, respondents from organizations that are currently using SSDs/Flash or have definitive plans to deploy SSDs or Flash, that 29% are using the SSDs for analytics applications.

### Simulation

In the recent survey, respondents from organizations that are currently using SSDs/Flash or have definitive plans to deploy SSDs or Flash, that 6% are using the SSDs for simulations.

### Transaction Processing

In the recent survey, respondents from organizations that are currently using SSDs/Flash or have definitive plans to deploy SSDs or Flash, that 28% are using the SSDs for OLTP applications.

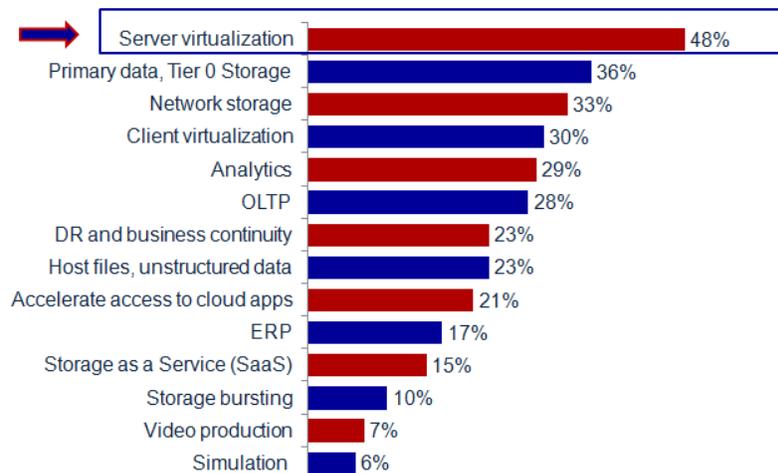
### VDI

In the recent survey, respondents from organizations that are currently using SSDs/Flash or have definitive plans to deploy SSDs or Flash, that 30% are using the SSDs for analytics applications.

### Server virtualization

In the recent survey, respondents from organizations that are currently using SSDs/Flash or have definitive plans to deploy SSDs or Flash, that 48% are using the SSDs for analytics applications.

#### If your firm currently uses/has definitive plans to deploy Flash/SSDs, how will it use it?



## Vendor/Product Profile

### Vendor Name: Dell

Product Name: Dell Storage SC Series (including Dell Compellent) and Dell Storage PS Series (including Dell EqualLogic) hybrid or all-flash array arrays, Dell PowerEdge Express Flash PCIe SSD, Dell PowerVault MD1220 DAS enclosure, Dell Storage Center Data Progression and Dell Fluid Cache for SAN

Product Type(s): All Flash Appliance/Array, Hybrid HDD and Flash Array and Tiering Software

Link to website: [www.dell.com](http://www.dell.com)

Dell has implemented SSDs/Flash capabilities in a variety form factors and software enhancements

- As hybrid or all-flash solutions in an SC Series/Compellent Storage Center SAN;
- As hybrid or all-flash solutions in the Dell EqualLogic PS Series Arrays;
- As the PowerEdge Express Flash PCIe SSD;
- As a drive alternative in its PowerVault MD1220 Direct Attached Storage enclosure; and,
- As Dell Fluid Cache for SAN and DAS software.

The Dell Compellent SC220 2U 24-drive enclosure features either all-Flash or hybrid-Flash storage. In the all-Flash configuration, all 24 drive bays are populated with a mix of six 400GB SAS-based SLC SSDs (good for write-intensive) and as much as 1.6TB SAS MLC SSDs. In the hybrid-Flash configuration, the drive slots can be filled with SSDs and as much as 1TB of 7.2K RPM SAS HDDs.

The Dell Compellent SC220 is managed by enhanced Storage Center 6.5 firmware that has been optimized for Dell's all-Flash and hybrid storage arrays. It provides the tiering of sub-LUN data across traditional HDDs, write-optimized SLC SSDs and read-optimized MLC SSDs.

The Dell Storage SC4020 all-in-one array is an all-Flash or hybrid storage array that offers 24 2.5" drive bays for 408TB capacity of 15K, 10K or 7.2K HDDs, write-intensive SLC SSDs and read-intensive MLC SSDs. It can connect to Dell's SC220 and SC200 drive enclosures.

Flash is enabled in two Dell EqualLogic arrays to meet a variety of customer needs: The all-Flash PS6210S and hybrid PS6210ES flash arrays contain 24 2.5" hot-pluggable SSDs and seven (7) SSDs and 17 2.5" 10K RPM SAS drives, respectively. Dell's EqualLogic line was one of the first in the industry to introduce SSDs into their systems and create hybrid arrays.

The PS Series/EqualLogic hybrid arrays auto-tier data on a continuous basis across SSDs and HDDs, writing and moving hot data to the SSD tier as needed.

The Dell PowerEdge Express Flash PCIe SSD is a storage device that contains SLC or MLC NAND Flash and is installed in a PowerEdge Server. The PCIe drive is a 2.5" form factor that is accessible from the front of the server, just like a

traditional drive. It features a capacity of 175GB or 350GB, as well as 1.6TB NVMe devices, and delivers performance in sequential read/write operations of as much as 1.8GBps. In random read/write operations, it features throughput of as much as 750,000 IOPs. The latency of the device is <50µ. As many as four PCIe SSDs can be installed in each PowerEdge Server.

The Dell PowerVault MD1220 Direct Attached Storage enclosure features as many as 24 2.5-inch SAS HDDs or 24 SAS SSDs.

In addition, Dell has the Fluid Cache for SAN acceleration software for its Compellent Storage Center SAN, which boosts the performance of the arrays. Previewed at Dell World in December 2013, Fluid Cache for SAN is a memory clustering technology acquired from RNA Networks that allows a shared memory pool of Flash storage to be created from the CPUs in the clusters and front-ends the disk capacity of the Compellent Storage Center SAN.

Dell Fluid Cache for SAN accelerates performance and application response time for application such as OLTP or virtual desktop infrastructure (VD). It uses Dell Express Flash PCIe SSD drives and Compellent data progression capabilities.

Management of Fluid Cache for SAN is from the Compellent Storage Center interface and includes configuration and deployment capabilities. Deployment of Fluid Cache for SAN requires Fluid Cache software, a minimum of three compatible Dell PowerEdge Servers supporting Dell PCIe Express Flash SSDs, connection to a 10/40GbE low latency private cache network and connection to a Dell Compellent SC 8000 controller and Dell Compellent SC200, SC220 and/or SC280 storage enclosures with rotating drives, flash or hybrid configurations.

Links to Data Sheets: <http://www.dell.com/us/business/p/dell-compellent-flash-optimized/pd>

<http://www.dell.com/us/business/p/equallogic-ps6210-series/pd>

[http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/Dell\\_PowerEdge\\_Express\\_Flash\\_PCIe\\_SSD\\_Spec\\_Sheet.pdf](http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/Dell_PowerEdge_Express_Flash_PCIe_SSD_Spec_Sheet.pdf)

<http://www.dell.com/downloads/global/products/pvaul/en/storage-powervault-md1220-specsheet.pdf>

<http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/DE5ENTC15917-DS1-Data-Prg-final.pdf>

<http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/dell-fluid-cache-for-san-specSheet.pdf>

<http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/fluid-cache-for-DAS.pdf>

## Signs of Maturity

A sure sign of market maturation is that the discussion has moved from shear feeds and speeds to the effects on application acceleration. Not all workloads are automatically improved by Flash, particularly sequential workloads in a non-virtualized environment. Rather than throwing out fear, uncertainty and doubt regarding all Flash and Hybrid Flash startups, the Tier 1 storage vendors have all developed their own all Flash and Hybrid Flash storage arrays. Here the dialogue has turned to enterprise storage features like automated replication and quality of service management at the workload level.

Accepted for several years as a 'Tier 0' of storage, Flash has now found its way into many Tier 1 applications, particularly in the web scale and cloud service providers. The financial benefit in these applications well offsets the cost per capacity, as more compute workloads can be handled on a single machine and converged compute and storage functionality on the same server becomes available to the mainstream service providers and enterprise datacenters.

The consolidation of many independent manufacturers by acquisition by foundry owners Toshiba, SanDisk, and Micron, as well as hard drive manufacturers Seagate and Western Digital have substantially changed the industry's landscape. With these consolidations, manufacturers are able to discuss solutions, which bridge from server-side Flash implementations to conventional SANs and NAS devices.

Interoperability between device types designed for Flash applications is increasing, with major manufacturers of related devices, like switches, host bus adapters and top of row devices taking into account the reduced latencies available from Flash devices, reducing bottlenecks in switching and protocols that were designed for much slower rotating media. Brocade's 'SSD Ready' initiative is an example. APIs are being developed for easy implementation of control and management of Flash devices including reporting of status regarding wear and potential of failure. Interoperability testing, at least for NVMe devices, which is the next wave of SSDs, has been institutionalized at the University of New Hampshire's Interoperability Laboratory.

Industry standard performance testing organization Storage Performance Council maintains a 'top ten' ranking of storage arrays in both total performance and price/performance. All devices in these top ten lists use solid state devices for primary storage. The following link is to the SPC site with the rankings:

[http://www.storageperformance.org/results/benchmark\\_results\\_spc1\\_top-ten](http://www.storageperformance.org/results/benchmark_results_spc1_top-ten)

## Our Take

The effects of Flash on the enterprise have been a whirlwind of deployment rates that consistently exceeded our projections. The advent of web-scale and hyper-converged infrastructures, which combine compute and storage on the same server, will only fuel this acceleration. We anticipate continued revenue growth from enterprises at near triple-digit rates.

The trends that we have been projecting, in terms of consolidation, will continue. Merger and acquisition activity will pare down the crop of current startups as well as the depletion of what has been viewed as a bottomless gold mine that has been providing venture capital to the startups.

Major manufacturers will continue to aggressively go after the Flash and hybrid array market as a way to protect and expand market share.

The NVMe devices will become an increasing factor in the market in the next year. These devices bring high performance with good standards for functionality that will enhance their usage. The devices are a perfect fit for the hyper-converged servers that provide compute and storage functionality in the same box.

The Flash foundries will produce volume products in the planar (flat) lithographies at 16/15 nanometers with MLC devices providing the largest capacity. TLC (three per cell) devices will begin to be used in consumer device applications with the help of signal processing algorithms.

Three-dimensional Flash will become generally available in about eighteen months from Samsung, Toshiba and SanDisk.

Alternative persistent memory device technology will not displace Flash for a decade or more.

Micron will become a player in the storage systems business within the next eighteen months, possibly through acquisitions.