



# Solid State Disks for HP StorageWorks Arrays Whitepaper

Abstract.....	2
What is Solid State Storage? .....	2
Forms of solid state storage.....	2
What technology is used in solid state storage? .....	3
Flash Memory Basics.....	4
NOR Flash.....	5
NAND Flash .....	5
Using NAND flash in solid state storage.....	6
Dialing in the desired storage properties .....	7
Where should SSDs be used? .....	8
What does the future hold?.....	10
Price trends .....	10
Competition.....	10
Applications .....	11
Standards .....	11
For more information.....	12

## Abstract

Magnetic disk drives have been in use for decades. Their behavior is well understood, predictable, and assumptions about disk performance are deeply ingrained in the industry and frequently in the applications that use disk storage. Magnetic disks are manufactured in huge quantities, and the industry has standardized on capacity points, performance ranges, reliability points, and price trends.

Solid state storage, especially flash memory based solid state storage, is quite new compared to magnetic disks. The explosion of the use of NAND flash memory in consumer products (cell phones, PDAs, cameras, media players, USB sticks) has created another high volume storage technology. But this technology is not backed by decades of experience, understanding, and standardized attributes, and will take years to reach such maturity. Therefore solid state storage, while already good and useful, will initially find use in the enterprise by those who can navigate these uncertainties and corresponding hype, such as early adopters and those with specific needs uniquely addressed by the technology. These users will naturally want to understand how this technology will affect performance, availability, reliability, and cost of their computer storage.

This paper will review types of solid state storage products, introduce the flash memory technology that underlies the recent explosion in solid state storage availability, and show how flash memory is used to build such a wide range of storage devices. Finally, it will explain the uses for which solid state disk technology is best suited in products like the HP EVA and XP arrays, and offer some opinions about the future of this technology.

## What is Solid State Storage?

We consider Solid State Storage to be any storage technology that:

- Has no macroscopic moving parts that wear out and limit performance.
- Retains data with no external power for a long time. In this paper, 'long time' means years, or something comparable to a disk on a shelf. (This excludes devices with internal batteries keeping data in volatile memory like DRAM).

## Forms of solid state storage

Solid State Storage may come in packages that have disk drive shapes and interfaces (Solid State Disks or SSDs) but can also come in non-disk forms.

Initially, SSDs will be the most popular deployment of solid state storage technology because disk drives are nice modular units with well defined, standard interfaces, and in an enterprise environment they are typically removable and hot-pluggable. Creating solid state storage in the form of a disk drive allows solid state technology to be most easily added to existing environments, and SSDs can then take advantage of the entire ecosystem that exists to manage disk storage. True to this trend, HP storage arrays initially offer solid state storage in the form of SSDs. XP supports 72 GB now, and will support 146 GB in March. EVA now supports 72 GB SSDs.

HP groups solid state disk devices into three classes: entry, mid-line, and enterprise. Entry storage is mainly intended for low-write uses such as boot devices, where product goals include power savings and smaller sizes, and where performance and wear life are not top considerations. Mid-line solid state storage is intended for high read performance and medium (but still much greater than magnetic disk) write performance, where wear out will take longer than the warranted life under specified write loads that are lower than the maximum possible sustained write loads. Enterprise solid state storage is intended for the most demanding uses, offering extreme read and write performance, and wear out that will take longer than the warranted life at the maximum sustained write performance, of course at the highest cost. As the market matures, the cost difference between entry and mid-line devices will

erode, causing entry devices to be marginalized. Some early consumer grade SSDs can also be found that have demonstrated write performance that ranges from low to abysmal, and a few even have significant wear and reliability problems; competition from more mature products will quickly drive these out of the market.

Solid state storage can also be found in external appliances that attach to servers using storage and network interfaces. In the past these have used DRAM storage, but are now also found in flash memory and hybrid DRAM + flash memory forms. These offer very high performance at tremendous cost, and so are affordable by few.

One penalty of building solid state storage that uses disk interfaces is that the device must implement the standard block storage protocols (For example, SCSI, SAS, and FC). Each of these protocols requires some processing to encode and decode. With magnetic storage, this processing overhead is miniscule because of the performance limits of the mechanical devices. With the very fast performance offered by solid state, this overhead becomes a much bigger consideration.

Already, some vendors are building block storage devices that connect directly to a PCIe interface, to get the highest possible performance by eliminating disk interconnect and protocol overhead. Other possibilities are being investigated by niche players, and surely more will be invented. Finally, contenders trying to replace NAND flash technology will continue to progress; several potentially disruptive technologies exist, and some might succeed.

## What technology is used in solid state storage?


Many mainstream and research technologies can be used in solid state storage:

- Solid state storage based on DRAM memory with optional batteries or external standby power and backup disks has been around for decades and offers some of the highest possible performance, though extremely high cost per bit has limited these to niche markets. Without backup power or batteries, the data is lost if power is lost.
- Devices based on flash memory are more recent. For the past several years, flash memory has offered a much lower cost per bit than DRAM, and requires no external power to retain data. Flash memory has become very well known due to its use in portable consumer devices such as cameras, media players, cell phones, and USB "thumb drives".
- Semiconductor-scale magnetic storage technologies, such as M-RAM or F-RAM, were developed to offer nonvolatile random access memory, but have never achieved the density or cost per bit of flash memory or even DRAM.
- Phase change memory promises behavior that improves with shrinking geometry, much faster write and erase times than flash memory, and much better wear out behavior, but also has not achieved the density or cost per bit of flash memory.

Because research and early technologies take years to develop and perfect, flash memory will remain the dominant technology used in solid state storage for at least the next four years, so let's take a closer look at it.

# Flash Memory Basics

A single flash memory storage cell is based on the same kind of transistor used in modern digital integrated circuits, but with the added twist of a floating gate that is not electrically connected to anything.




Due to the marvelous and peculiar nature of quantum physics, and because of the extreme thinness of the dielectric layer beneath the floating gate, appropriate voltage levels across the floating gate can energize electrons enough that they 'appear' in the floating gate and get trapped (writing or programming), or move out of the floating gate (erasing). A trapped charge affects the threshold at which a voltage on the control gate can turn the transistor on or off (reading). Reading does not require electrons to move across the dielectric layer, so reads happen at full integrated circuit logic speeds. But writes and erases require a certain number electrons to move across the dielectric barrier so the write and erase voltages must be applied long enough for this to happen. Because of this, writes and erases are much slower than reads in flash memory technology.

Unfortunately, the movement of the energetic electrons across the dielectric layer slowly damages this layer. Eventually the damages reaches a point where the floating gate can no longer hold a charge, or the electrons cannot move across the barrier at the designed rate. This causes the phenomenon called wear. Flash memories can only support a certain number of write and erase cycles before the cells 'wear out.' When a cell wears out, it can no longer be used to reliably store data. The number of times a cell can be programmed and erased before it wears out is referred to as endurance.

Another important concept is retention, which measures how long a flash memory can retain data while turned off. For some uses of solid state storage, it is important that data be retained at least as long as well understood technology such as disk drives. Interestingly, for flash memory wear and retention are related. As a cell wears out, the length of time it can correctly retain data is shortened.

In the original flash memory designs, a single cell represents one bit of data. An erased cell represents a '1' state, and a cell programmed with the appropriate charge trapped in the floating gate represents a '0' state. This is called single level cell (SLC) flash memory. In the relentless march to improve densities and reduce cost, engineers found that they could store multiple charge levels in the floating gate. Today, commercially available parts can store four different charge levels in one cell, allowing the cell to store two bits of data. This is called multi level cell (MLC) flash memory.



Most consumer devices now use MLC flash technology to get the lowest cost per capacity. Unfortunately, MLC devices wear out much faster than SLC devices. This paper will further explore this difference below.

Up to this point, we've looked at flash memory from the viewpoint of a single cell. But a flash memory chip is a collection of many cells, and there are two fairly different ways of organizing the cells into a flash memory chip called NOR and NAND, along with a few hybrid schemes from some vendors that mix properties of both.

## NOR Flash

NOR flash was the original organization of cells in flash memory, and remains the preferred technology for certain uses such as boot code or embedded device code storage. NOR memory allows complete random access read, so a processor can run boot code right out of NOR memory. NOR memory is generally designed so that reads are reliable without any error correction, because often NOR flash is used to run boot code before any error correction hardware or software had been set up. NOR flash memory was in fact originally developed to replace code storage ROMs so that the code that would have been stored in ROM could be updated a limited number of times. While NOR flash can be read with true random access, it must be written and erased in larger units.

## NAND Flash

Because of the support for true random access reads, NOR flash memory requires a lot of internal wiring to allow random access to the individual memory words. This wiring takes up valuable space on the chip. NAND flash was invented to increase storage density at the expense of random access. NAND flash memory must be read and written in pages, and erased in blocks that are multiples of pages, so NAND flash is naturally a block storage technology. In exchange for the loss of random reads, NAND flash offers a lower cost per bit for bulk storage than NOR flash. Because NAND flash is always read in pages, NAND flash is generally designed and specified assuming some amount of page error correction to function properly. This allows parts with a few faulty bits to work correctly, further increasing chip yields and reducing costs. Because of these characteristics, solid state storage devices primarily use NAND flash memory.

Now we shall briefly look at how NAND flash memory is used. A modern high capacity NAND flash might have 4 KB pages for read and write, and an erase block size of 64 pages (256 KB). A 4 KB page is actually 4096 bytes plus a few hundred bytes of extra storage for use by the flash memory controller for ECC storage, metadata storage such as counting erase cycles for wear tracking, bad page management, and so on. A 1 GB memory would have 256K pages and 4K erase blocks.

To read a page out of flash memory, the controller identifies the page to read, the memory chip reads the entire page into an on-chip buffer, and then the buffer can be streamed out of the chip. The page read is very quick (30 uS), and the streaming can support quite high data rates (>20 MB/sec per chip).

To write a page into an erased block, the page data is streamed into a buffer, and then the flash chip writes the entire page from the buffer to a page of NAND flash cells in parallel. The streaming also supports the high data rates similar to the reads, but the page programming time is much slower than the page read time (200–800 uS). The time required to erase a block is even slower than the programming time (1–2 mS).

The wear life of a NAND flash part varies by process technology and flash memory type. Typically, SLC NAND flash blocks are specified to wear out after 100,000 erase cycles, using 1 bit error correction per page. MLC NAND flash blocks wear out after 5,000–10,000 erase cycles, using 4 bit error correction. If SLC is used with 4 bit or higher error correction, the 100,000 cycles specified could be exceeded, though the amount of improvement varies greatly by manufacturer and process technology, so care must be taken not to assume too much improvement without actual testing.

One final and less well known peculiarity of flash memory is the phenomenon of read disturbs. In a usage model of infrequent writes and intensive reads, a page may be written once and read tens of thousands of times or more. The read signals, while not as energetic as the write signals, do affect the electrons trapped in the floating gate. Enough reads can disturb the charges, causing some bit errors, but this will be detected and handled by the error correcting logic, which can trigger the controller to rewrite the corrected data, avoiding problems. Unfortunately, the read signals can also disturb the stored charges in adjacent pages, which is not so readily detected without taking extra steps to catch these problems.

## Using NAND flash in solid state storage

Now that we've examined the behavior and peculiarities of NAND flash memory, how do vendors use these devices to build a solid state storage device, such as a solid state disk, with desirable properties? In one sense, the flash memory is the media of a solid state disk. As we briefly covered above, NAND flash has peculiar behaviors and limitations that must be managed. These are very different than the limitations of magnetic storage, but no less important.

The first and most important limitation is the wear out problem. It is easy to imagine a data record that is updated very frequently. If the data record could be updated in-place in a flash memory device, then a MLC device with a wear out of 5000 program-erase cycles would wear out the page containing that data in a matter of seconds. Even with a SLC device, the page containing the hot data would wear out in a minute or so. Flash memory controllers solve this problem by tracking the amount of wear in each block and moving data around to distribute the wear. There are many algorithms for doing so, but the end result is that the wear is spread over all the erase blocks in the device. When blocks wear out, they are marked as bad and not reused. Devices are configured with spare capacity, but eventually a sufficient number of blocks will wear out that the spare capacity is exhausted and the device becomes unusable.

The organization of NAND flash memory also introduces write performance issues. If every write required reading the contents of a whole block, erasing a new block, and writing the block back, then even for a small write the write performance of the device would be abysmal. Some consumer devices actually exhibit this behavior, which can be seen by write performance as low as 10 writes per second, which is terrible even by magnetic disk drive standards.

A better approach would be to have a pool of blocks already erased and waiting for incoming writes. With this approach, the performance is limited only by the page write time, and the block erase is not a factor. This requires some extra capacity to have erased blocks ready for writes, and it also requires the some background processing to do the housekeeping necessary to always have

some erased blocks ready to use. Extra capacity has the side benefit of increasing the overall wear capacity of the device. It also requires the ability to map individual writes or whole pages to available erased space, and the ability to keep these complex maps correctly.

If a write of any size causes a whole page to be moved to relocate the written data, this ratio between the actual write and the page size is called write amplification by some vendors. This is used to help determine how fast the device wears out with different write mixes. Another approach would be to use something akin to a log-based file system where writes are appended to the end of an erased block, but then the wear overhead comes from managing complex maps and garbage collecting free space. In either case, the actual wear is a fairly complex function of the mix of write sizes and the internal algorithms used to manage space and wear, and the latter are generally proprietary secrets of the vendor that builds the solid state storage device.

As you can imagine, any solution that offers decent performance and wear behavior must maintain a fairly complex map of user data to flash storage locations. The management of this map is itself another complicating factor. Solid state storage devices contain some DRAM which is used for several purposes, including:

- Buffering of data from the external interfaces
- Maintaining mapping data
- Caching data

Some vendors use short term power storage such as super capacitors to flush the mapping data and any cached data from DRAM to flash in the event of a power loss to the device. Other vendors rebuild the mapping data from flash memory structures in the case of unexpected power loss, but this rebuild can take up to several minutes. If the DRAM holds write-back cached data, some vendors choose to simply lose the cached data, similar to what is done in magnetic disk drive caches, while others do not use a DRAM cache or guarantee that cached data is flushed to flash. HP is requiring vendors that supply our enterprise SSDs to build devices that do not have long rebuild times, and do not have write back caches that lose data.

## Dialing in the desired storage properties

NAND flash memory of both the SLC and MLC types has been used to build solid state disks that exhibit a huge range of performance and reliability. This section briefly looks at how such a range of results is possible using the same underlying technology.

The following design choices are employed to get a desirable solid state storage device:

- Number of channels of flash memory: the solid state storage controller can use many channels to concurrently access many flash memory chips. Enterprise SSDs may use as many as 20 separate channels of flash memory chips and can drive these channels concurrently, making extremely high performance devices possible. Mid-line SSDs use at least 4–8 channels of flash memory.
- Capacity over-provisioning: extra capacity means extra cost, extra wear life and extra space, which is kept erased and ready for writes in order to provide the highest performance. There is a lot of variation between vendors, but mid-line devices generally have on the order of 10 percent capacity over provisioning to get performance, and enterprise device may have up to 100 percent over provisioning to get long wear life and extreme performance. Clearly this is a big factor in the higher cost of enterprise devices, and this is a cost differential that will not go away with time.
- Amount and use of DRAM for mapping structures, data cache, buffers, and scratch pad uses.
- Choice of SLC for highest performance and longest wear, or MLC for lowest cost per capacity at the cost of performance and wear life. Today enterprise-class solid state storage uses SLC flash, entry devices primarily use MLC (except for some premium branded flash memory devices), and mid-line server and storage devices can be found with both MLC and SLC, trending more towards

MLC over time. As capacities increase, there is more space in which to spread the wear, and this makes MLC SSDs last longer. Soon some high capacity MLC devices will be able to sustain maximum performance random writes for the warranted life of the device, but cannot do so for sequential writes.

- Proprietary mixes of algorithms for managing wear, write allocation, data mapping, erase-ahead, garbage collection, and so on, all of which affect performance and wear life under different types of workloads and write size mixes.

As we can see above, solid state storage can be very complex, but a properly designed device can offer excellent performance and wear properties. Because most solid state devices today exist in the form of solid state disks, let's look at some example of what to expect. Enterprise-class SSDs can offer extraordinary random small block random read and write performance (50,000 read/sec, 15,000 writes/sec), and latencies are equally impressive (as low as 100 microseconds). These can cost more than ten times as much as an enterprise grade 15K RPM magnetic disk on a \$/GB basis, but offer as much as 100 times the random I/O performance. Sequential performance of SSDs is also good, but no where near as huge a step up, because magnetic disks have quite good sequential performance. An enterprise grade SSD can have three times the sequential performance of an enterprise magnetic disk, but at more than ten times the cost this is not a compelling use for SSDs unless some other attribute such as the very low access latency is worth the cost.

Regarding wear out behavior, enterprise grade SSDs, which are being used in HP StorageWorks EVA and XP arrays, are designed to run at maximum possible write speed for the designed warranty life of the device (typically 5 years to match enterprise magnetic disks). Because mid-line devices offer somewhat lower performance and will wear out ahead of their warranted life if driven at maximum write traffic, HP enterprise arrays use enterprise grade SSDs.

Because of the limitations of flash chip density and chip packaging solutions, and because of cost, SSDs do not offer the same capacity range as similarly sized magnetic disks. At current technology points, the maximum capacity offered by a SSD in a given disk form factor (2.5" or 3.5") will be roughly the same as the minimum capacity offered by magnetic disks in that same form factor.

## Where should SSDs be used?

Up to this point, this paper has focused on the performance benefits of solid state storage. While this is certainly the most important benefit in the enterprise space, solid state storage has other useful attributes. These include:

- Lower power—Solid state storage, especially entry and mid-line disks, use less power per device than magnetic disks. This can be very useful to reduce power and cooling needs, or to extend the operating time of battery powered devices such as notebook computers.
- Far lower power per I/O—Because of the extremely high I/O rates supported by solid state storage, the amount of power consumed per I/O is vastly lower than that of magnetic disks. Enterprise applications that deploy lots of magnetic disk 'spindles' to get a high random I/O rate would see great power savings by using SSDs, though clearly this would also result in much less total capacity. If this can match an applications needs, then solid state storage offers a great benefit.
- Noise—Because there are no moving parts, solid state storage is silent.
- Shock and vibration—Magnetic disks with heads flying microscopic distances from a spinning disk surface can suffer from head crashes from the shock of being dropped even small distances while operating. SSDs do not have this problem, again making them very suited for mobile devices such as notebooks. Immunity to vibration also makes solid state devices suitable for harsh vibration environments such as use in vehicles of all sorts.

- Environmental tolerance—Solid state storage can operate at a wider temperature range and altitude range than magnetic disks, making them the only viable solution for some extreme environments.
- Malleable—Magnetic disks require a rigid spinning platter in a sealed protective case, which limits the possible mechanical shape of disks. Solid state storage device simply use electronic components on circuit boards, allowing them to take a wider variety of shapes for specialized uses. The SSD designed for the HP BL495c virtualization blade shows offers an example of this use.

For enterprise uses, the most valuable benefit of solid state storage devices is the very impressive random access performance and very low latencies. They also offer very good sequential performance, but not good enough to justify the increased cost compared to magnetic disk drives. Example of uses where the excellent random performance and low latency help include transaction processing, data mining, write intensive random access workloads such as databases and other data stores (Exchange servers, for example). Because of the high cost per bit, applications where the most performance sensitive data can be segregated from bulk data that has lower performance demands are best suited to using solid state storage as a higher performance tier of storage. In cases where customers find it cost effective to purchase more expensive 15K RPM enterprise grade disks and then short-stroke them over a fraction of their total capacity to get higher performance, enterprise grade SSDs can offer many times the performance at a cost per capacity that is close to that of magnetic drives.

Because of the extremely high random performance of enterprise SSDs, a modern array designed around the performance of large farm of magnetic disks will encounter internal bottlenecks when using a much small number of SSDs. Attempts to use more SSDs than the array controller can support will waste the performance of the additional SSDs and the cost paid for this performance. Therefore each array has guidelines for the maximum number of supported SSDs and how they should be configured.

Because the cost per bit of SSDs is much higher than magnetic disk, SSDs are not cost effective for bulk data storage for all classes of data. Clearly their deployment should be focused where the added cost can be justified, such as the most performance sensitive workloads, or where other benefits justify the increased cost.

Because solid state devices have no moving parts, some have claimed that this eliminates the need for data redundancy such as that offered by array controllers or application redundancy. This is simply not the case. While a solid state device does not suffer the classes of mechanical failures that magnetic disks have, there are still plenty of other failures cases, including component failures, solder and connection failures, and firmware failures. Generally, enterprise grade SSDs have an MTBF specification that is two times better than magnetic disks or less, and there isn't a particularly long track record from most vendors to back up this data with real experience. If such a device fails, the flash chips may still contain the customer data, but it inaccessible because the overall device has failed. For customers that require high availability and reliability, external forms of redundancy are as needed as ever.

One unique attribute of flash-based solid state storage is the wear out. Most types of failures in magnetic disks and electrical components are random failures that are modeled statistically as MTBF. Because of the error correction and wear management algorithms in such solid state devices, the devices can actually know how much capacity is worn out or near end of life, and can report that information back to the application or hosting system. This gives new opportunities for measuring the useful life of a device and scheduling replacement of devices nearing the end of the wear life.

# What does the future hold?

## Price trends

Some in the storage industry have claimed that solid state disks will completely replace magnetic disk drives. This will not happen any time soon; there will be a price per capacity advantage for magnetic disk drives at some storage tier for a long time. This claim is usually made assuming a continued 50 percent per year price drop in flash, but ignoring the historical price trends of magnetic disk drives. The following chart shows the price per GB trends of DRAM, flash memory chips, and disk drives.



If disk price per capacity returns to its long term 50 percent historical trend, that would match the flash trend, in which case flash solid state storage devices would never reach disk price per capacity. On the other hand, if disks stay at their more recent 30 percent trend, then the lines would cross, but far enough in the future that the more likely scenario is another disruptive technology will come along and change the landscape before flash can replace all uses of magnetic disk storage.

Initially customers that are willing to short stroke the most expensive 15K RPM magnetic drives to get the highest possible random access performance will benefit greatly from enterprise SSDs at the least additional cost. As SSD costs come down, they will be cost effective for an increasingly wide range of customers, but again predominantly for random access workloads. Because magnetic drives are quite good at sequential workloads, the benefit of SSDs in those uses is much less for their high cost.

Because flash memory is a semiconductor technology, it will undergo the same process shrinks and resulting cost improvements that the industry as a whole experiences. But in the case of flash memory, smaller geometries actually increase the wear out problem, so the bright minds in the industry have some work to do to avoid these problems.

## Competition

The solid state storage business has become very busy recently. Small companies and startups that have pioneered the flash solid state disk business are getting a lot of attention. The major disk drive vendors are moving into this area, because they see the competition affecting their higher performance drive business. Flash memory vendors also want to enter this business to move up the

integration ladder, and to smooth out a business model that today is subject to the wild swings of consumer device demand. The various vendors all have their favorite approaches to solving the challenges of using flash memory, with corresponding patents and other intellectual property. As the best approaches to building SSDs emerge from this competition, the surviving competitors will likely converge on these approaches, just as has happened in the disk drive business over several decades.

## Applications

As is the case with the wider IT industry, storage arrays will initially take advantage of solid state storage technology in the form of SSDs. HP offers enterprise grade SSDs for use in the XP and EVA arrays.

Based on the attributes of enterprise SSDs described above, current enterprise SSDs offer much higher random read and write performance and much lower latencies than magnetic disks, at a higher price per capacity. SSDs will therefore be most useful in environments where a subset of data or metadata would benefit greatly from very high IOPS and very low latency. As competition drives SSD costs down, the use cases where SSD costs are justified broaden. For a while, the relative cost per capacity of enterprise SSDs compared to magnetic disk will get better as increased competition brings SSD costs in line with the actual cost to manufacture the devices in volume. In the longer run, both magnetic disk and SSD cost per capacity will continue to decline, so magnetic disks will retain a price per capacity advantage. However the absolute cost for an SSD will continue to decrease, increasing the set of customers for which SSDs have sufficient value to be worth using.

In StorageWorks arrays, SSDs and magnetic disks cannot be mixed within a redundancy group (Disk Groups on EVA, Parity Groups on XP). Therefore SSDs are best grouped to form a higher performance tier of storage, with appropriate data placed on this faster, more expensive tier. HP will be working with key application vendors to help characterize the benefit and recommend best configurations to use SSDs with these applications. Because SSD availability on HP storage arrays is so new, this type of information is not yet available. Look for additional publications from HP to provide this information as it becomes available.

Most storage-intensive applications have been designed for years based on the behavior of magnetic disk storage. As SSDs become more affordable and popular, the applications will change internally to take better advantage of the benefits of solid state storage. This will take time, but will increase the benefits of using SSDs.

## Standards

The flash memory business is already adopting standards to help with multi-vendor interoperability. An example of this is the ONFI standard for the interface to the basic flash chips. NVMe proposes to standardize a nonvolatile memory controller interface.

In flash storage devices, standards are needed for ways to measure performance, wear endurance, wear reporting, and so on. HP is actively participating in the standards bodies that are trying to address these problems.

## For more information

Resource	Web address
HP on Solid State Storage Technology	<a href="http://h71028.www7.hp.com/enterprise/cache/609040-0-0-121.html?jumpid=go/solidstate">http://h71028.www7.hp.com/enterprise/cache/609040-0-0-121.html?jumpid=go/solidstate</a>
SNIA whitepaper: Solid State Storage 101, An introduction to Solid State Storage	<a href="http://www.snia.org/apps/group_public/download.php/35796/SSSI%20Whi%20Paper%20Final.pdf">http://www.snia.org/apps/group_public/download.php/35796/SSSI%20Whi%20Paper%20Final.pdf</a>

© 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Intel, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

4AA2-4509ENW, February, 2009

