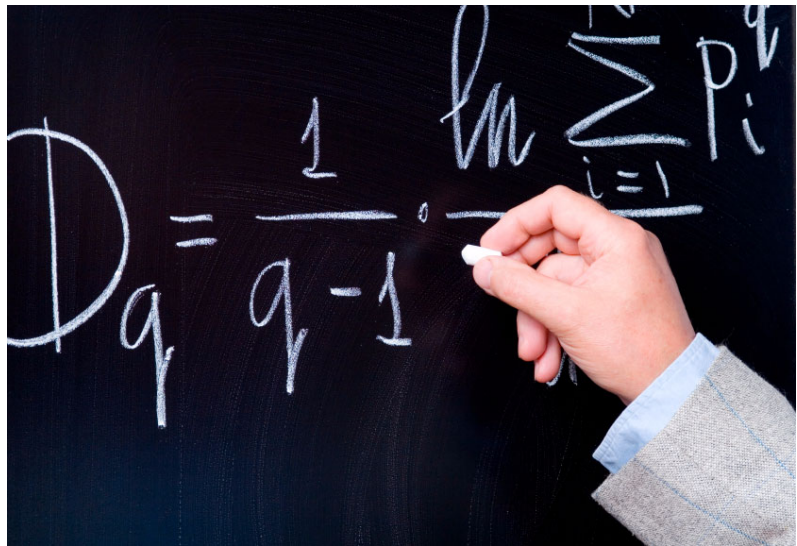


## WHITE PAPER

# NAND EVOLUTION AND ITS EFFECTS ON SOLID-STATE DRIVE (SSD) USEABLE LIFE



## EXECUTIVE SUMMARY

Solid-state drives (SSDs) are attractive alternatives to rotating hard disk drives (HDDs) because they eliminate the single largest failure mechanism of HDDs — moving parts. The concern relative to NAND flash-based SSDs is that there is a limitation in the number of write/erase cycles (endurance) and therefore the SSD may not meet long-term system deployment requirements, especially in 24/7 applications.

OEMs want to answer the question, "How long will this SSD last?" in some time unit like hours or years. They want the parameters of that answer to be calculated similarly between vendors so apples-to-apples comparisons can be made. This is difficult if not impossible to accomplish because the two most significant variables in the endurance calculation are usage model and write amplification. While OEMs may be able to monitor their usage model in terms of data written from their system to the controller on the SSD, they will not know how the SSD controller manages NAND flash on the back end (this is the write amplification factor). There could be a significant difference between the amount of data the OEM thinks he is writing and the amount of data actually being written to the NAND.

This white paper examines the basics of NAND flash and some algorithms SSD vendors use to manage it. This document also proposes an endurance metric that separates the parameters the SSD manufacturers control (type of NAND used, write performance, and write amplification) from those the system OEM can control (capacity and write duty cycle). This parameter, called LifeEst, yields a value of "write years per gigabyte." Multiplying this value by the capacity in gigabytes and dividing by the write duty cycle yields a lifetime estimate in years.

Even with well-modeled applications, calculations are at best theoretical. A better methodology, one that yields real world results, is to use a tool within the application itself to monitor the exact wear of the NAND flash and report that data back to the host system. SiliconSystems patent-pending SiSMART technology is one such tool.

Simply plug SiliconSystems' SiliconDrive SSD into the application and run it for a reasonable test period. At the end of the test period, read the usage data from the SiliconDrive and make a simple extrapolation to useful life. As an example, say an OEM runs a SiliconDrive for a week and reads the SiSMART data as 0.2% used. Assuming the application does not change, this drive should last for 500 weeks. This lifetime estimation is simple to calculate and yields the most accurate value because it is based on real world, application-specific usage data. This tool can also be integrated into the OEM's software and used during system deployment to monitor endurance and make system-level decisions to eliminate down time due to excessive wear.

## TABLE OF CONTENTS

Introduction .....	1
NAND Flash Basics .....	2
NAND Trends.....	4
SSD Reliability .....	7
Equations, Calculations, and Measurements .....	9
Conclusion .....	13

## INTRODUCTION

It is well known that NAND flash, the primary media in solid-state drives (SSDs) is changing at a rate never before seen in semiconductor technology. NAND manufacturers are driven by the unrelenting quest for lower cost per bit to accommodate the needs of next generation consumer devices from cell phone handsets to music and video players to low cost and ultra-mobile PCs. This lower cost per bit is opening up new applications for SSDs and is making them more compelling when compared to rotating hard disk drive (HDD) alternatives.

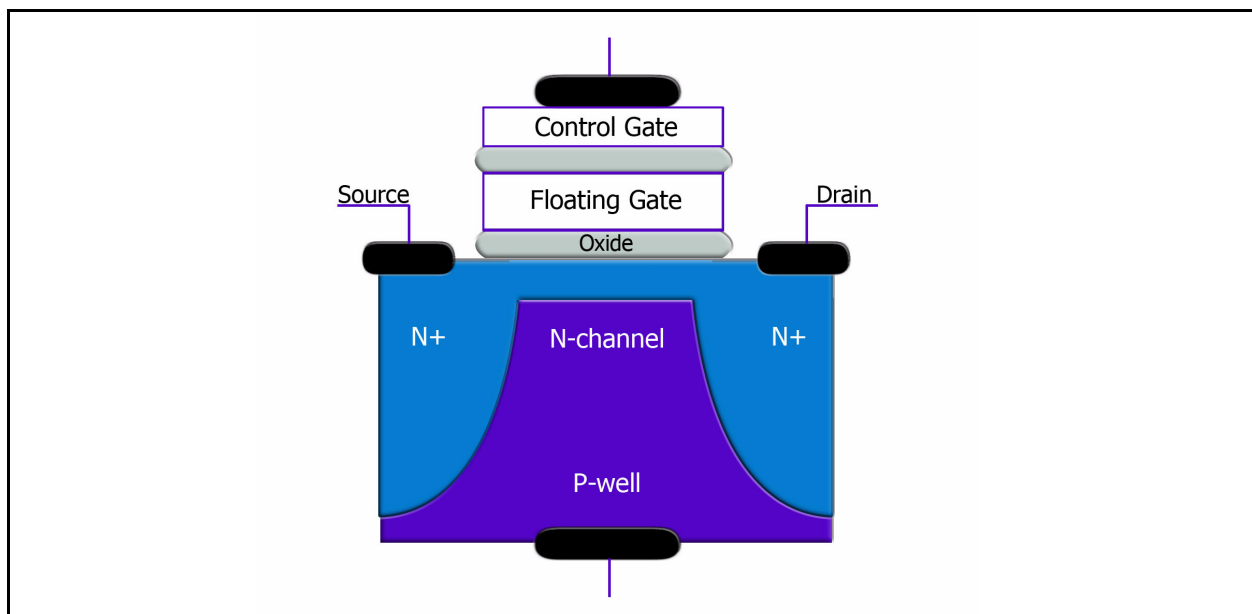
NAND vendors have been very aggressive in their efforts to store more bits of data per NAND component at a lower cost per bit and they have attacked this issue in two ways. First, they seek to reduce the manufacturing process geometry, which shrinks the size and therefore the cost of the NAND cells (transistors) that actually store the data. The smaller the cells, the more of them can be etched on a given area of a semiconductor wafer. Process geometry shrinks, plus some additional solid-state physics issues not discussed in this paper, results in larger densities per NAND component. Second, NAND vendors offer components that can store multiple bits of data per cell in order to double, triple and soon quadruple the number of bits of data that can be stored per component.

The trade-off for smaller process geometries and lower cost per bit is in component reliability. NAND flash components in their most common TSOP packages are getting bigger, faster and cheaper, but not "better." Better in this case is defined as "more reliable" in terms of endurance (the number of program/erase cycles) and data retention.

This does not mean that NAND-based SSDs are inherently unreliable. In fact, at the raw media level, NAND flash is more reliable than the raw magnetic media used in HDDs. It simply means that SSD controller technology faces the same issues as the HDD controllers that preceded them — how to take advantage of the lower cost per bit while maintaining acceptable reliability levels for the usage model at hand. SSDs have the luxury of attacking this issue without being burdened with what has traditionally been considered the biggest reliability headache — the mechanics of rotating media.

### NAND FLASH BASICS

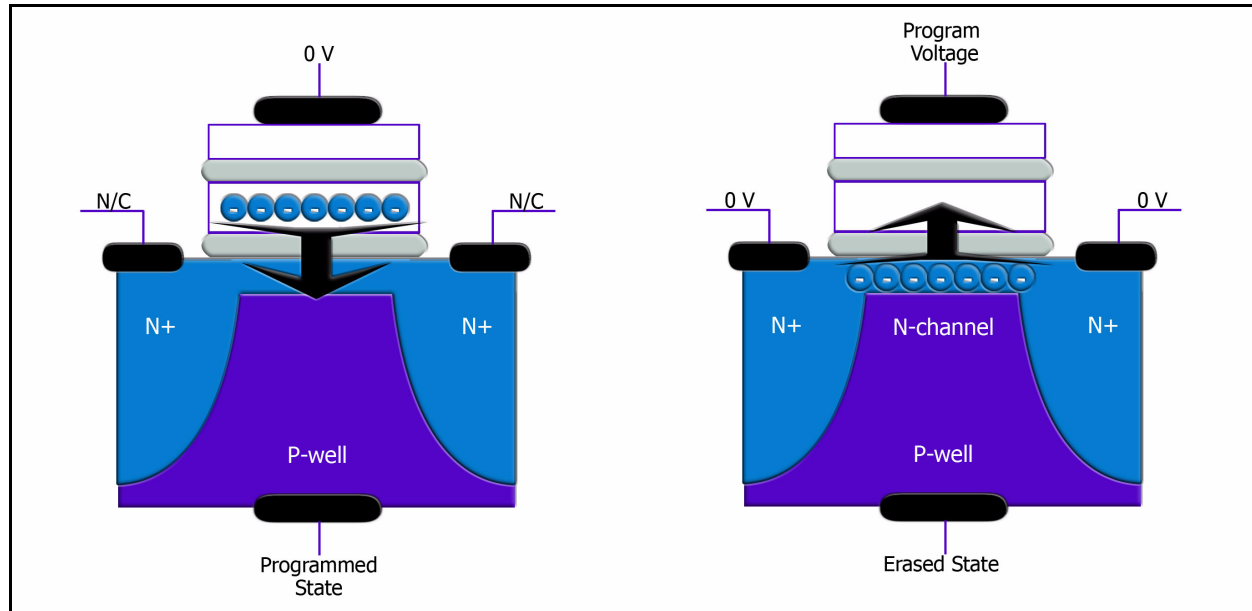
Since NAND flash is the most commonly used non-volatile solid-state media for SSDs, it is helpful to understand a few fundamentals of NAND flash physics. The following figure illustrates the basic structure of a NAND flash cell.



**Figure 1: NAND Flash Cell**

The cell is really just a floating gate transistor. The transistor works by forming a conductive channel between the source and drain. A voltage applied to the control gate results in the formation of a negatively-charged channel which conducts electricity from the source to the drain. When the control voltage is removed, the channel disappears and no conduction takes place.

The cell also contains a secondary floating gate which is not electrically connected to the rest of the transistor. The voltage at the control gate required for channel formation, called the voltage threshold, can be changed by modifying the charge stored on the floating gate. A "programmed" state is defined as charge added to the floating gate and an "erased" state is defined as no charge on the floating gate as illustrated in [Figure 2 on page 3](#).



**Figure 2: NAND Cells in Programmed and Erased States**

The more electrons that can be stored on a floating gate, the larger the difference in voltage thresholds between a programmed state and an erased state. As process geometries shrink, the cells get smaller and fewer electrons are stored on the floating gate. This means that there is less of a voltage threshold difference between programmed and erased states and therefore a smaller margin of error if electrons "leak back into the substrate." The net effect is that while a smaller process geometry means a lower cost per bit, it comes at the expense of long-term reliability as voltage margins decrease.

In addition to shrinking process geometries, NAND manufacturers offer components that can store multiple bits per cell. Single-level cell (SLC) NAND (traditionally called binary) stores one bit per cell. It uses one voltage threshold to distinguish between logical 0 and 1. Multi-level cell (MLC) NAND has traditionally been associated with storing two bits per cell, but recent technology advances have allowed three bits per cell to be stored. Two-bit-per-cell (2X) MLC requires four different voltage levels to uniquely identify the proper bit combinations. Three-bit-per-cell (3X) MLC requires eight different voltage levels and n-bit per cell NAND requires  $2^n$  different levels. This is illustrated in [Figure 3 on page 4](#).

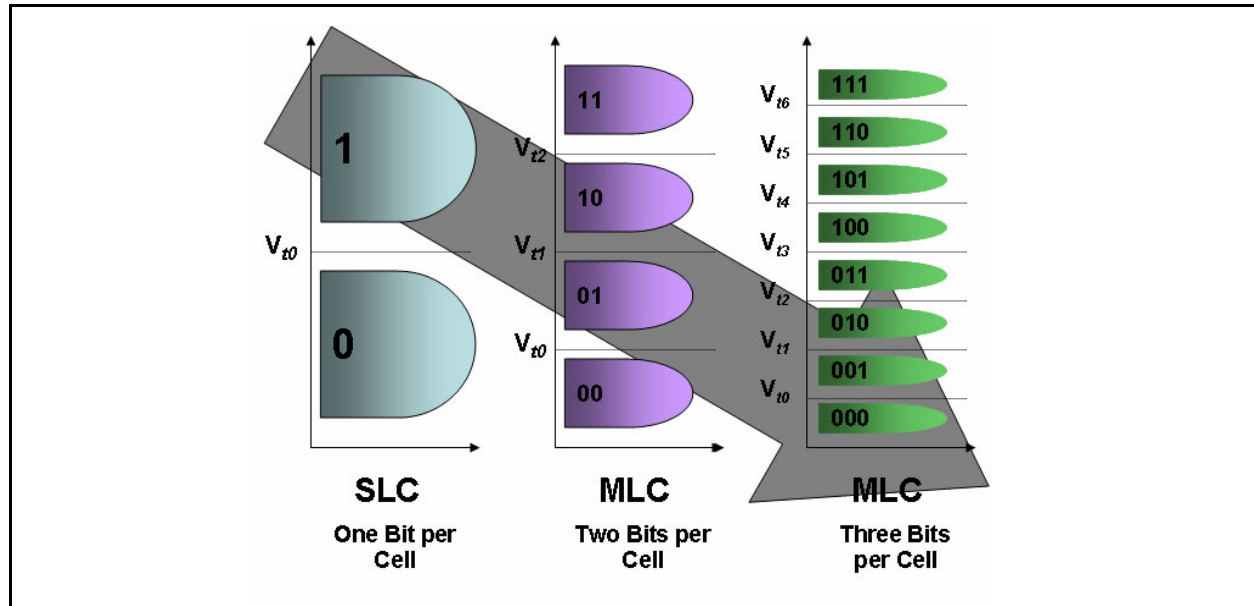


Figure 3: Reliability vs. Bits per Cell

MLC also has a write performance trade-off against SLC (read speeds are generally the same for both technologies). MLC takes longer to program than SLC because the charge state of the floating gate must be carefully monitored during programming since there are multiple voltage thresholds. The more bits per cell, the slower the programming must be to ensure the proper floating gate charge state. This results in slower program time and therefore slower write speed for MLC.

## NAND TRENDS

Now that basic NAND technology has been defined, it is useful to review the current trends in the NAND industry to better understand the challenges facing SSD manufacturers as they integrate these components that are getting bigger, faster and cheaper, but not better.

### BIGGER

NAND component densities per TSOP (the most common IC package) are doubling at a rate that exceeds Moore's law. SLC components have reached 64 gigabits (Gb) using four internally stacked 16Gb monolithic die. That density can be doubled again by using a type of TSOP stacking where the X-Y dimensions stay the same while the Z height increases slightly. Alternatively, new packaging techniques are allowing eight die to be stacked in some types of BGA packages. MLC technology doubles these densities and have enabled SSD vendors to achieve capacities of up to 512 gigabytes (GB) moving to 1 terabyte (TB) in traditional 2.5" form factors — significantly closing (and in some cases eliminating) the "capacity gap" between SSDs and HDDs.

### FASTER

During the manufacturing process, memory cells are etched onto a common section of the NAND substrate forming a single block of memory called an erase block. This is the smallest segment of a NAND component that can be erased at one time. Depending on component density, block size can range from 16 kilobytes (KB) to 512KB — moving quickly to 1 megabyte (MB).

Programming the memory cells is performed one word at a time (i.e., cell by cell) and usually an entire page (e.g., 4096 bytes) is programmed in a single operation. Depending on component density, page size can be 2KB or 4KB (increasing to 8KB in 2009 for 16Gb monolithic SLC and 32Gb monolithic MLC). Page size and block size are fundamental specifications for NAND components.

NAND flash vendors have been able to find ways to erase and program these blocks and pages in the same amount of time even as their densities double. This effectively increases the write speed by enabling more bits to be programmed in the same amount of time. In addition, high speed interfaces like Open NAND Flash Interface (ONFi) and other alternatives being discussed in the JEDEC standards group are speeding up the communication links between the NAND and the controller, further increasing speed.

### CHEAPER

Specific NAND component manufacturing costs and prices will not be explored in this white paper. Regardless, it is extremely important to understand that major changes have taken place in the NAND market in order to understand how the past will (and will not) relate to the future.

In 2003, SLC NAND dominated fab output as the largest producer, Samsung, controlled 60% of the market and exclusively manufactured SLC. Toshiba, the second leading supplier with approximately 30% market share, split its output between SLC and 2X MLC. The final 10% of the market was split between several vendors, including Renesas, Fujitsu and others (source: [Web-Feet Research](#)).

In 2008, there are four major NAND suppliers. Samsung still leads the market with 40% share. FlashVision, a joint venture that supplies NAND to SanDisk and Toshiba is number two, followed by Hynix and IM Flash, which produces NAND for Intel and Micron. The output mix has changed dramatically over the last five years with roughly 90% or more of today's NAND output being 2X MLC and less than 10% being SLC (source: [iSupply](#)).

When originally introduced, 2X MLC was roughly half the cost of the equivalent density SLC component. The explosion of consumer electronic devices (most notably the iPhone and iPod from Apple) that have chosen 2X MLC over HDDs, combined with the significant competition from new entrants into the market, have driven 2X MLC component prices down significantly faster than their SLC counterparts — to the point where now the technologies are virtually uncoupled from a price perspective.

To achieve some historical perspective, 8GB SLC-based SSDs sold for approximately \$4,000 in 2004. Today, 16GB MLC-based SSDs routinely sell for less than \$200 — double the capacity for 1/20 the price in just four years.

### NOT BETTER

"Better" in this case is defined as more reliable in terms of endurance and data retention — the length of time a charge remains on the floating gate after the last program cycle. Multiple writes and erases to a NAND flash cell breaks down the oxide barrier over time, making it more difficult to "keep the electrons in place." Therefore, there is an inverse relationship between endurance and data retention — the more writes and erases, the shorter the data retention.

Endurance and data retention must always be specified together in order to be meaningful. A higher endurance can be specified if the data retention is low. Conversely, a lower endurance specification can result in a higher data retention value. The exact nature of this relationship varies between components and is not linear, but it is becoming increasingly more important to understand as NAND manufacturers continue to sacrifice endurance and data retention for lower cost per bit.

JEDEC (JESD22-A117A, JESD47) specifies the relationship between endurance and data retention this way:

100% rated endurance	1-year data retention
10% rated endurance	10-year data retention

There are varying relationships between these two endpoints. For example, in a disk cache application, multi-year (or even multi-month) data retention may not be required since by its very nature, data in a disk cache is considered transient. In this instance, a higher endurance may be specified while specifying low data retention.

NAND flash manufacturers must make value judgments and engineering trade-offs between cost per bit and endurance. The most well-known trade-off is in MLC vs. SLC, where a 2:1 or 3:1 cost benefit is achieved for a 10:1 reduction in rated endurance. Process geometry shrinks also result in smaller and less expensive NAND at the price of fewer electrons being stored on the floating gate of the transistor as discussed earlier.



The following table shows typical NAND flash endurance ratings at the component level:

**Table 1: Typical NAND Flash Component Endurance Ratings**

Process Geometry	Rated Endurance with 1 Year Data Retention
5xnm and larger SLC NAND:	100K cycles with 1-bit ECC
5xnm 2X MLC NAND:	10K with 4-bit ECC
4xnm and 3xnm SLC NAND:	100K with 4-bit ECC
4xnm and 3xnm 2X MLC NAND:	5K with 8-bit (or more) ECC

## SSD RELIABILITY

SSD controller companies, working in conjunction with NAND manufacturers, are employing several techniques to extend the life of an SSD at the drive or system level. Techniques such as wear-leveling, error detection and correction, and block management algorithms all contribute to increased system-level endurance for a given level of data retention.

### WEAR-LEVELING

Wear-leveling is commonplace in the industry today, but there are several different methodologies employed. The most robust wear-leveling algorithms wear data evenly across the entire SSD. While some vendors call this static wear-leveling because any static data (operating system files, executable programs, look-up tables, and so on) will be moved to different physical block addresses. Others call it dynamic simply because it "sounds better." This paper will call wear-leveling over the entire SSD, independent of host operating and file systems, "full wear-leveling."

SSD vendors targeting consumer electronics and personal computing applications such as netbooks or notebooks understand that the usage models for these applications involve significant human interaction. As a result, these vendors may only wear-level over a small portion of the SSD. Performing this type of "partial wear-leveling" relies on the host operating system (OS) to do its own wear-leveling (disk defragmentation for HDD is a result of this phenomenon). Relying on the host to provide some of the necessary storage management algorithms enables consumer-based SSD vendors to use a cheaper controller that is suitable for human interaction usage models but which comes up well short of meeting the 24/7 operating requirements of many embedded and enterprise systems.

Wear-leveling efficiency is a relatively new parameter being defined by many SSD vendors. This parameter reflects the maximum deviation of the most-worn block to the least worn block over time. It is usually discussed in terms of a percentage. For example, a wear-leveling efficiency of 3% indicates that there is only a 3% differential of wear on the most worn block as compared to the least worn block.

### ERROR CORRECTION

While challenging to implement, the concept of error correction is relatively simple: more is better. However, there are some practical limitations of how many bits can be corrected. First, there are engineering trade-offs between number of bits corrected, controller cost and complexity and overall SSD performance. Next, the number of bits per sector that can be corrected is limited by how many bytes the NAND component puts aside for ECC data. Smaller density NAND components may only allow for 4-bit correction. NAND manufactured at newer process geometries set aside enough bytes to store 16 bits of ECC data per 512 byte sector.

### STORAGE MANAGEMENT ALGORITHMS AND WRITE AMPLIFICATION

NAND flash is considered "mostly good memory," which must be proactively managed. Bad blocks are flagged in a table and are managed by the SSD controller. Spare blocks are used in wear-leveling and other write/erase operations to increase the endurance at a system level. In general, the number of spares in an SSD is 1 to 2%, but it can be as high as 50% in high reliability applications. This is usually accomplished by over-provisioning the SSD (i.e. providing additional NAND capacity specifically to address reliability issues.)

Write amplification is a measure of the efficiency of the SSD controller. Write amplification defines the number of writes the controller makes to the NAND for every write from the host system. The concept stems from the fundamental mismatch between erase block sizes (256KB for 50nm SLC), page sizes (4KB for 50nm SLC) and sector sizes (512 bytes). Long, continuous writes map over this mismatch, but most embedded/enterprise applications do not stream data. Instead, they transfer data in a series of shorter, more random transactions.

The minimum write size from an SSD controller to the NAND is usually the page size — in the above example 4K. Most SSDs on the market must erase before writing so a 4K write from the host will, at worst case, require a whole erase block (256KB) to be erased and written. To state it another way, 256KB are written from the SSD controller to the NAND for a 4KB write from the host to the SSD controller. In this case there is a 256:4 or 64:1 write amplification. In truth, write amplification is somewhere between perfect (1:1) and worst case which is defined as (erase block size/page size:1).

## EQUATIONS, CALCULATIONS, AND MEASUREMENTS

Many SSD vendors are struggling to classify endurance in terms that are meaningful to OEMs and end users. “So many write/erase cycles per logical block” may be useful as a "checklist" specification, but it does not answer the real question, "How long will the SSD last in my application?" OEMs need to understand SSD life in terms of time — years, months, days — instead of "cycles." Defining and measuring the usage model is critical to making this translation.

In the commercial/consumer space, SanDisk is touting the concept of Longterm Data Endurance or LDE. LDE is based on [Business Applications Performance Corporation \(BAPCo\)](#) human interface usage models for professionals, students and personal users. It defines the total number of data writes allowed in the SSD based on one year of data retention. It is specified in terabytes written (TBW) and is directly related to the capacity of the drive.

The LDE value for an SSD is measured by limiting the capacity of the drive and running it until it fails. Assuming linear usage, the LDE value can be specified for the drive. Below is an example (this data is for illustrative purposes only and does not necessarily reflect LDE values of any product):

LDE Rating for 64GB SSD: 40TBW — a 32GB SSD composed of the same technology would have an LDE rating of 20TBW.

BAPCo Usage Model: Professional writes 80GB per week

Lifespan: 40TBW/80GB per week = 512 weeks = just under 10 years

While this specification is useful for netbook and similar applications, there are a couple of limitations. First, it is really only targeted for human interaction usage models and does not address the needs of 24/7 embedded or enterprise systems. Second, this model does not take into account the randomness of data and the associated write amplification factors.

This white paper offers a more targeted approach for enterprise system OEMs by reviewing two generic methodologies based on 24/7 usage models with a requirement for one year data retention.

The first is a streaming application where the host system writes data from logical block address zero (LBA0) to LBA<sub>n</sub>. In this case, the write performance of the SSD is maximized and the effects of wear-leveling and block management are minimized. The lifetime, measured in years, is defined by the following equation:

$$\text{Lifetime} = \frac{\text{Endurance Rating} * \text{GB of Storage} * (.0325)}{\text{Maximum Write Speed (MBps)} * \text{Duty Cycle}}$$

- Endurance Rating — The block level endurance that has been traditionally specified as 100K, 10K, or 5K. Use the value "5" for 5K, "10" for 10K, and so on. Many vendors do not give out this information because the NAND is changing so rapidly. Consequently, many users "try out" different values and adjust capacities accordingly.
- 0.0325 — Constant derived from "endurance rating in thousands of cycles," "KB-to-GB," and "seconds-to-years" unit conversion
- Duty cycle — The percentage of write cycles to (read cycles + idle time)

**Example:** A video surveillance application monitors public transportation. If no "event" occurs, the drive over-writes the data. A 64GB drive capable of a sustained data rate 32MBps based on MLC NAND rated at 5,000 write/erase cycles per block yields:

$$\text{Lifetime} = \frac{(5) * 64 * 0.0325}{32 * (\text{Duty Cycle})} = \frac{0.32 \text{ Years (3.9 Months)}}{(\text{Duty Cycle})}$$

Very few embedded applications stream data. Most are database or transactional applications. In these IOPS-type applications, the numerator stays the same with the exception of a 1024 KB to MB conversion factor, but the denominator is significantly different. IOPS measurements are made using the industry-standard IOMeter benchmark. In IOMeter, the user can define file size and the percentage of reads and writes. The write IOPS rating is the output of IOMeter based on the desired file size.

The concept of write amplification comes into play because simply monitoring the host writes (IOPS rating) does not yield the proper information. Duty cycle must also be considered.

$$\text{Lifetime} = \frac{\text{Endurance Rating} * \text{GB of Storage} * (.0325) * 1024}{\text{Write IOPS Rating} * \text{File Size in KB} * \text{Write Amplification} * \text{Duty Cycle}}$$

**Note:** The equation above is considered "worst case" based upon 100% random data.

- Endurance Rating — The block level endurance that has been traditionally specified as 100K, 10K, or 5K. Use the value "5" for 5K, "10" for 10K, and so on. Many vendors do not give out this information because the NAND is changing so rapidly. Consequently, many users "try out" different values and adjust capacities accordingly.
- 0.0325 — Constant derived from "endurance rating in thousands of cycles," "KB-to-GB," and "seconds-to-years" unit conversion
- 1,024 — Constant for KB to MB conversion
- IOPS rating — Number of write input/output per second as measured by IOMeter
- File Size — The file size at which the IOPS rating is measured

- Write amplification — The number of writes at the NAND level for each host write. This value is related to usage model, but the worst case, for 100% random writes, is a value of 64. This value is based on the ratio of NAND erase block size to page size. If the file size is larger than the page size, then the worst-case write amplification is erase block size divided by file size.
- Duty cycle — The percentage of write cycles to (read cycles + idle time)

**Example:** A voicemail system manufacturer is considering a 32GB SSD to replace a rotating disk drive. The drive uses SLC NAND that is rated at 100K endurance with 200 write IOPS for an 8KB file. The drive does not specify a write amplification factor so a value of 32 (256KB block/8KB file) will be used. The OEM estimates the write duty cycle at 25%.

$$\text{Lifetime} = \frac{(100) * 32 * (.0325) * 1024}{200 * 8 * 32 * 25\%} = 8.3 \text{ Years}$$

It can be clearly seen that there are direct correlations between capacity and useful life. Assuming a full wear-leveling scheme, doubling the capacity doubles the useful life. This may not be the case for SSDs that implement a partial wear-leveling scheme.

An alternate way to consider this model is to figure out the capacity required based on a product deployment requirement and specified usage model.

**Example:** A network virtualization appliance requires a one year product deployment based on a usage model of 3,000 write IOPS on a 4KB file at 50% write duty cycle. The preferred drive vendor rates the SSD at 300K endurance with 3 months data retention, but does not specify the write amplification factor. How large does the SSD need to be?

$$\text{Lifetime} = 1 \text{ year} = \frac{300 * \text{GB of Storage} * (0.03247) * 1024}{3,000 * 4 * 64 * .5}$$

$$\text{GB of Storage} = 39 \text{ GB}$$

A 32GB SSD may not do the job but a 64GB product will provide a nice safety margin.

The real useful life of an SSD is governed by three parameters: SSD technology, capacity and usage model. OEMs can use capacity and usage model to determine their useful life based on the SSD technology. To that end, SiliconSystems proposes a new metric to measure SSD technology — LifeEst. LifeEst measures the SSD technology by specifying number of "write years per GB" the SSD can achieve and is defined by the following equation:

$$\begin{aligned} \text{LifeEst} &= \frac{\text{NAND Endurance Rating} * \text{Unit Conversion Constant (UCC)}}{\text{Write IOPS rating} * \text{File Size in KB} * \text{Write Amplification}} \\ &= \text{Write years per GB} \end{aligned}$$

The UCC is calculated by:

$$\begin{aligned} \text{UCC} &= \frac{[1,000 \text{ (Endurance Cycles)}] * [1,048,576 \text{ (KB/GB)}]}{[3,600 \text{ (sec/hr)}] * [24 \text{ (hr/day)}] * [365 \text{ (day/yr)}]} \\ &= 33.25 \end{aligned}$$

In the preceding example, LifeEst can be specified by:

$$\text{LifeEst} = \frac{100 * 33.25}{200 * 8 * 32} = .065 \text{ Write Years per GB}$$

Useful life can now be calculated easily from the following equation:

$$\text{Lifetime (Yr)} = \frac{\text{LifeEst} * \text{Capacity (GB)}}{\text{Write Duty Cycle (\%)}}$$

A 32GB drive would last about 2.1 years at a 100% write duty cycle. A 25% write duty cycle yields about 8.3 years of useful life.

LifeEst is the first endurance parameter that is independent of capacity and usage model. It specifically measures performance and storage management algorithms implemented in the SSD. Using NAND components with a higher endurance rating will increase LifeEst. Decreasing the write amplification will also have a positive impact on useful life, as will slowing down the SSD.

### MEASUREMENTS

Equations are great, but measurements are better. The above discussion is based on theory. A better model is to use an endurance monitoring function like SiliconSystems' patent-pending SiSMART technology which is integrated into all SiliconDrive Products. Using this methodology, the SiliconDrive will yield real-world, application-specific usage information without requiring software development on the part of the OEM. Simply load the SiliconDrive into the application and run it for an hour, a day, a week, a month, or any time that makes sense. SiSMART reports the usage percentage during that time.

Upon test completion, simply take the SiliconDrive out of the test system and plug it in to any Windows or Linux based computer and run the SiliconSystems SiliconDrive Utility to measure the usage at the NAND media level.

In this case, the Lifetime calculation is very straightforward:

$$\text{Lifetime} = \frac{\text{Test Length}}{\text{SiSMART Usage}} \times \frac{\text{Years}}{\text{Test Length}} \times 100$$

**Example:** A point-of-sale OEM is considering an SSD to replace an HDD. The OEM runs a 16GB SiliconDrive for one week and uses the SiliconDrive utility to determine that the drive was 0.25% used.

$$\begin{aligned} \text{Lifetime} &= \frac{1 \text{ Week}}{0.25} \times \frac{1 \text{ Year}}{52 \text{ Weeks}} \times 100 \\ &= 7.69 \text{ Years} \end{aligned}$$

Again, in a SiliconDrive that performs full wear-leveling, doubling the capacity will double the useful life.

For a particular SiliconDrive part number and capacity combination, the SiSMART usage data percentage will remain constant for the specific usage model — independent of block-level endurance of the specific NAND component used. SiliconSystems patented SolidStor storage management technology shields enterprise system OEMs from the myriad NAND changes and ensures consistent performance over the SSD life.

## CONCLUSION

Relative to SSD endurance, the simple question remains, how long will the SSD last *in my application*? OEM system designers have traditionally approached sizing their storage needs by first starting with the premise of so many cycles at block level. They then look at the size of the OS and program files, the amount of data to be collected and they come up with an SSD capacity.

This methodology was adequate when NAND component technology was not changing rapidly and the predominant NAND component in use was SLC NAND that was capable of 24/7 usage. Today, NAND is changing almost as rapidly as the qualification cycles of some OEMs and those OEMs need to take a more system-level approach. They first need to determine how long their product must be deployed in the field. Next, they need to determine the usage model (and just as importantly how to measure and predict it). From there they can specify the proper capacity for the required SSD technology. Finally, the OEM must partner with a storage vendor that has a business and technology philosophy in place that keeps form, fit, function and reliability constant throughout the product lifecycle.

Download the SiliconSystems' SiliconDrive Utility, which includes the patent-pending SiSMART utility, at the SiliconSystems Partner Portal: <http://www.siliconsystems.com>.

## **CONTACT US**

### **Worldwide Headquarters Office:**

SiliconSystems, Inc.  
26840 Aliso Viejo Parkway  
Aliso Viejo, CA 92656

Tel: 949.900.9400

Fax: 949.900.9500

Website: [www.siliconsystems.com](http://www.siliconsystems.com)

For technical questions, contact at: [appeng@siliconsystems.com](mailto:appeng@siliconsystems.com).

For Marketing/Sales, contact: [info@siliconsystems.com](mailto:info@siliconsystems.com).

SiliconSystems' performance tests, ratings, and product specifications are measured using specific computer systems and/or components and reflect the approximate performance of SiliconSystems' products as measured by those tests. Any difference in system hardware or software design or configuration, as well as system use, may affect actual test results, ratings, and product specifications. SiliconSystems welcomes user comments and reserves the right to revise this document and/or make updates to product specifications, products, or programs described without notice at any time. SiliconSystems makes no representations or warranties regarding this document. The names of actual companies and products mentioned herein are the trademarks of their respective owners.

SiliconSystems®, SiliconDrive®, SiliconDrive II®, SiSecure™, PowerArmor™, SiSMART™, SiKey™, SiZone™, SiProtect™, SiSweep™, SiPurge™, SiScrub™, SiliconDrive USB Blade™, and the SiliconSystems logo are trademarks or registered trademarks of SiliconSystems, Inc. and may be used publicly only with the permission of SiliconSystems and require proper acknowledgement. Other listed names and brands are trademarks or registered trademarks of their respective owners.

© Copyright 2008 by SiliconSystems, Inc. All rights reserved. No part of this publication may be reproduced without the prior written consent of SiliconSystems.