

Temperature Considerations for Industrial Embedded SSDs

Introduction

NAND flash-based SSDs and memory cards continue to be the dominant storage media for most industrial-embedded systems. Historically, industrial-embedded systems have had limited capacity needs, ranging from several megabytes (MB) to a few gigabytes (GB). Their functions have been well defined, they have required real-time operating systems with a relatively small storage footprint, and the data being stored has been relatively small. In most industrial-embedded applications, system designers can still fit SLC-based SSDs into their storage budget so they can get all the benefits of longer product lifecycles, high endurance and wide temperature range operation. MLC does not provide much of a price break for 8GB and below storage solutions, especially when balanced with the shorter lifecycles and endurance concerns. When viewed from a total cost of ownership perspective, SLC makes the most sense.

But as embedded systems continue to add enhanced functionality and deal with “bigger” data, capacity requirements can skyrocket to tens or hundreds of gigabytes. At higher capacity points, NAND cost dwarfs the cost of other SSD components and storage budgets come under pressure. OEMs can no longer afford high-end SLC SSDs and must look at less expensive alternatives that still give them the capacity required by their system design. They must trade-off down-the-road requalification and endurance considerations to get to market with an affordable solution.

This white paper discusses the affects that extended high temperatures have on SSDs. It discusses endurance and data retention characteristics of SSDs with two different types of NAND flash media – 2xnm SLC (Virtium PE class) and 1znm MLC used in standard (Virtium CE class) and iMLC (Virtium XE class) modes.

Background

“How long will this SSD last in my application and how often will I need to requalify?” Those are the two “How long will this SSD last in my application and how often will I need to requalify?” Those are the two questions industrial embedded system OEMs ask most frequently. The requalification question is fairly straightforward. MLC will need to be requalified every one to two years. SLC should be available for three to five years and depending on the type of SLC NAND used, however, it could be a lot longer than that. The “how long will it last in my application” is a much more difficult question. SLC offers a significant reliability advantage over MLC and will be offered here as a baseline, but this paper will include the endurance and data retention characteristics of MLC-based product for those higher capacity applications where SLC is not economically viable.

In order to choose the right SSD for an application, it is important to understand reliability factors and their relationships. Industry standards body JEDEC has defined two application classes, one for client (personal) computing and one for enterprise (multi-user) applications. There is no “industry standard” for embedded computing since applications are so diverse, but the JEDEC models provide a good framework to use as a guide.

For each application class, JEDEC defines a data pattern or workload, an acceptable error rate, operating and data retention temperatures, and an acceptable period of time the data must be retained in the power-off state. Many SSD manufacturers, including Virtium, use the JEDEC enterprise application class in product datasheets to address endurance, but it is important to note that endurance and data retention can change dramatically based on workload and operating temperature. Use these datasheet values as a guide to compare SSDs, not as an absolute. Since most application workloads and temperatures differ greatly from the JEDEC enterprise application class, expected endurance and data retention will similarly vary.

Class	Workload	Active Use (power on)	Retention Use ¹ (power off)	Functional Failure Requirement	UBER Requirement
JEDEC Client	JESD219 Client	40°C, 8 hours per day	30°C for one year	≤ 3%	<1 sector in 10 ¹⁵ bits read
JEDEC Enterprise	JESD219 Enterprise	55°C, 24 hours per day	40°C for three months	≤ 3%	<1 sector in 10 ¹⁶ bits read

1: After endurance requirement has been met

Table 1: JEDEC Application Classes

Detailed explanations of the terms mentioned can be found in JEDEC JESD218 and JESD219 documents. For the purposes of this paper, workload, UBER, and functional failure requirement will be held constant while examining the effects of endurance, high temperature and NAND configuration on power-off data retention

Definitions

1. Drive writes (DW) – the number writes to the full capacity of the SSD. Many SSDs datasheets specify Terabytes Written (TBW). Since TBW is related to user capacity, DW is just a normalized factor for any capacity for a given workload. DW is defined by the following equation:

Drive Write (DW) =	NAND P/E Cycles
	Write Amplification (WA)

$$\text{Drive Write (DW)} = (\text{NAND P/E Cycles}) / (\text{Write Amplification (WA)})$$

2. NAND P/E Cycles – the number of program/erase cycles supported by NAND, based on a given error correction (ECC) capability of the SSD controller, at an assumed temperature (40°C, 104°F) and an assumed data retention requirement (1 year for SLC, MLC and iMLC).

3. Write amplification (WA) – the amount of data written to NAND versus amount of data coming over the interface. This value is greater than 1, assuming no data compression, and is the result of a mismatch in the page size of the NAND (the minimum write unit) and the block size of the NAND (the minimum erase unit). Write amplification is determined by firmware algorithms and workload. It can range from an ideal state of 1, where data comes across the interface in large file, sequential transfers and is aligned to NAND flash page boundaries; to a worst case of block size divided by page size. In general, the more random the write

workload, the higher the write amplification. The subsequent graphs are based on a mixed industrial workload which results in a write amplification of $WA = 4$.

Drive Write Example

An SSD using 1nm MLC NAND flash rated at 3,000 program/erase cycles at 40°C (104°F) with one year of data retention is subject to a workload that results in a write amplification factor of 4. The total drive write value (also shown in Figure 5 below) is:

Drive Write (DW) =	3000	= 750 Total Drive Writes
	(4) x (1)	

Drive WSo if the product deployment requirement was for five years, the number of drive writes per day would be:

A 30GB SSD would handle 12GB per day for five years while a 60GB SSD would accommodate 24GB per day. Understanding how much data needs to be written for what period of time will be a significant aid in properly sizing the SSD capacity.

Data Retention

There are two types of data retention – power-on and power-off. Power-on data retention for most SSDs, including Virtium’s StorFly® and TuffDrive® products, is virtually unlimited. This is because most newer, high-end SSDs implement patrol read and patrol scrub algorithms where the SSD firmware periodically reads all LBAs and repairs or refreshes them as necessary. For the rest of this paper, data retention is assumed to be power-off data retention, representing systems that may be sitting on a shelf prior to deployment or after being decommissioned.

Before continuing, a brief physics discussion is in order. Figure 1 shows the structure of a NAND flash cell.




Figure 1: Structure of a NAND Flash Shell

The data value is determined by the number of bits per cell and the voltage level read by the SSD controller. The voltage level is determined by the number of electrons on the floating gate of the transistor. Over time, electrons on the floating gate can leak through the oxide layer back to the substrate. The more electrons leak, the more the voltage changes and the higher the chance of a bit error. Too many bit errors - more than the SSD controller can correct - results in uncorrectable, and eventually system, errors.

So the stronger the oxide layer, the better the data retention. Oxide strength is determined by two factors –

endurance and temperature. The more program / erase cycles, the weaker the oxide layer becomes. In terms of temperature, think of the oxide layer as ice on a lake. When programming, electrons get injected from the substrate onto the floating gate. The colder the temperature, the more difficult it is to program. The hotter the temperature, the easier it is to program. The converse is also true. The colder the temperature, the more difficult it is for electrons to leak back into the substrate. The hotter the temperature, the more leakage can occur.

So the best-case scenario is to program at higher temperatures and store at lower - this is reflected in the JEDEC application classes above. Use this information as a rule of thumb rather than a guide, because the oxide layer can heal itself with dwell time between writes and temperature adjustments.

Effects of Drive Writes and High Temperature on NAND Flash-based SSDs

It is now time to examine the effects of drive writes and temperature on the power-off data retention for various configurations of SSD. As discussed in the introduction, comparisons will be drawn between the same capacity SLC, MLC and iMLC drives at various drive write (DW) points at various temperatures.

Early Stage Data Retention

Many OEMs are concerned with initial or early stage power-off data retention. This is because they may configure and test their system at their manufacturing facility and then it could sit weeks or even months on the shelf prior to being deployed. They are concerned that firmware, operating systems and other configuration data could be lost before the system is deployed.

Figure 2 shows data retention characteristics for the three different SSD configurations at drive writes $DW \leq 25$. If the system is stored in a relatively benign environment, data retention should not be an issue. Remember, however, that these values are assumed constant. It would seem as though few storage facilities would sit at 85°C (185°F) for over two months.




Figure 2: Power-Off Data Retention Days Prior to or Early Deployment

Continuous Operation at High Temperatures and Data Retention

The worst-case scenario is continuously operating and storing at high temperatures. Figure 3 and 4 show the SSD classes at 70°C (158°F) and 85°C (185°F), respectively. Note that the graphs are in logarithmic scale since the data retention degrades so quickly, especially for MLC.

Figure 3: Power-Off Data Retention at 70°C (158°F) per Number of Drive Writes




Figure 4 – Power-Off Data Retention at 85°C (185°F) per Number of Drive Writes

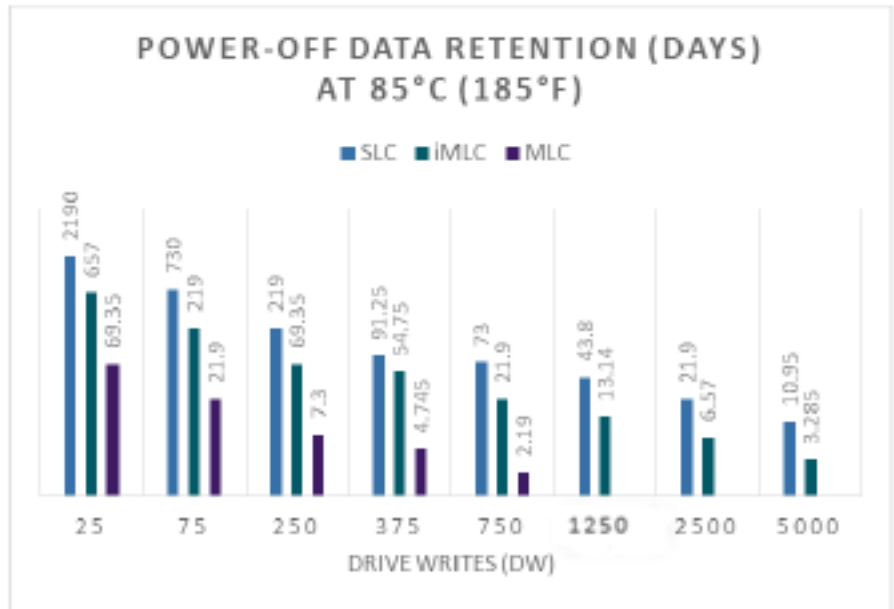
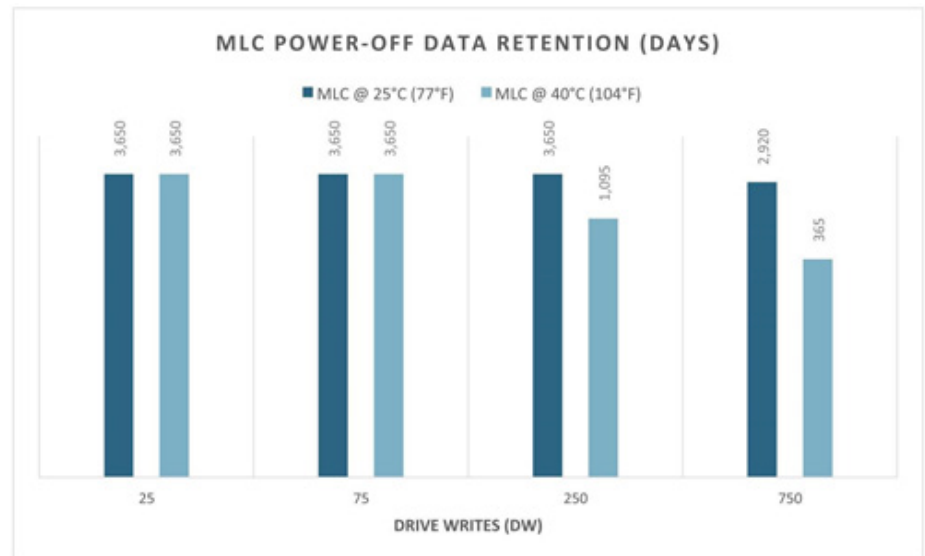


Figure 5 – MLC Power-Off Data Retention at 25°C (77°F) and 40°C (104°F) per Number of Drive Writes



The data shows the dramatic effects that temperature has on data retention for given workloads. For the same 750 full drive writes (0.4 drive writes per day for five years), SSDs operated and stored at 85°C (185°F) will only have two days of data retention, where as those drives at 40°C (104°F) will have one year and those at room temperature 25°C (77°F) will exhibit characteristics of nearly eight years of data retention.

Conclusion

The drive for higher and higher capacity points is forcing some embedded systems OEMs to consider using MLC-based SSDs instead of traditional SLC SSDs to achieve storage budget targets. While SLC SSDs are faster, more reliable and don't need to be requalified nearly as often, if capacities get too large, the price premium for SLC may become burdensome and customers need to use MLC.

Once that decision is made, it is important to understand workload, operating temperatures and data retention requirements – in particular, how long does the SSD need to retain data when power is removed from the system? When SSDs are relatively new, data retention is not a real concern. As SSDs are deployed for a longer period of time, especially at higher temperatures, OEMs need to size the SSD to give them enough retention time to take care of the data in the event the system loses power and needs to be brought back up.

Proper SSD sizing is the key. Start with capacity needed for operating system, lookup tables and user data. Then try to determine service life, approximate drive writes per day and operating temperature.

Please keep in mind that when investigating I-temp SSDs, endurance and data retention characteristics change with workload and temperature. Industrial temperature SSDs are guaranteed to operate over the given temperature range and metrics like performance, power and response to commands are consistent throughout. That does not hold true, however, for data retention at a given endurance value, so it is important to plan workload, capacity and data back-up accordingly. With SLC, it is not a big concern. With MLC, care must be taken.

Virtium manufactures memory and storage solutions for the world's top industrial embedded OEMs. For two decades we have designed, built and supported our products in the USA - fortified by a network of global locations. Our world-class technology and unsurpassed support provide a superior customer experience that continuously results in better industrial embedded products for our increasingly interconnected world.

© Copyright 2016. All rights reserved. Virtium®, vtView®, Storfly® and TuffDrive® are registered trademarks and vtGuard™ a trademark of Virtium LLC. All other non-Virtium product names are trademarks of their respective companies.



30052 Tomas | Rancho Santa Margarita, CA 92688
Phone: 949-888-2444 | Fax: 949-888-2445
www.virtium.com