

swissbit®

# Selecting flash storage for industrial applications



## Learning to Choose

# Selecting flash storage for industrial applications

---

*"The aging of NAND chips, a NAND flash-inherent effect, is a key influence on parameters like write- and read performance or retention."*

---

Selecting flash storage modules for embedded systems and industrial applications is a task that should not be taken too lightly. Some requirements are obvious like the temperature and vibration resistance specification defined by the application. Others require a strategic TCO-approach (Total Cost of Ownership): when comparing product life cycles of, for example, a medical device or a production machine with life cycles of IT equipment, it becomes apparent that it is vital to have a supplier who guarantees long-term availability of modules that were meticulously qualified for a particular ap-

plication. But there are further aspects that determine the selection of a memory module for embedded industrial applications. This white paper explains the aging effects of NAND chips and the counter measures the manufacturers of flash memory take, to expand endurance, retention, security and performance. Understanding the internal mechanisms of an SD Memory Card or an SSD is a prerequisite to ask the right questions and make a well-founded selection of storage devices matching the requirements of the application.

## Introduction

There is a wide choice of flash memory technologies available on the market today including TLC, MLC, SLC, pSLC and 3D NAND. Without an understanding of the basic mechanisms of storage and the effects at storage level, selecting the most suitable flash memory products for particular applications is difficult. The first step is knowing which questions to ask providers. Depending on the application, different aspects need to be considered when selecting memory modules for embedded industrial applications:

- read and write speed
- endurance (the lifespan of flash media)
- retention (the lifespan of stored data)
- data security in case of power failure
- temperature and vibration resistance
- long-term availability of the choice of product

---

The aging of the NAND chips, which is an effect specific to flash, plays a significant role in many of these aspects.

## Aging of NAND-Cells – Shifted threshold voltage

*"The main reason why SLC is still one of the preferred solutions for industrial applications, despite the higher cost, is because of the higher endurance of 100.000 P/E cycles compared to 3.000 for MLC."*

Typically, only a limited number of block erase cycles are possible for the cells of NAND flash devices. Each time the programming voltage is applied and the tunnel effect is generated, a strong electrical field accelerates electrons in the direction of the tunnel oxide. Some of these "hot electrons" receive

enough energy by the programming voltage so that they become trapped in the gate oxide instead of tunneling onto the floating gate. Consequently, the threshold voltage changes over time and ultimately the cell is no longer readable. (Figure 1)




Figure 1

Multiple programming of the cell injects hot electrons in gate oxide instead of floating gate. These electrons are trapped temporarily and cause a VT shift of this cell. Read fails will increase over time. Block erase cannot recover the fails, damage is permanent, and the block needs to be marked as "Bad block".

## Conductive paths

The formation of conductive paths in the oxide layer is another effect of aging. It causes the cell's charge to gradually decrease and with it, the hold bit. (Figure 2) Exposure to high temperatures speeds up the effect. It is particularly prevalent in cases where the number of applied P/E (program/erase) cycles approaches the wear out limit, causing

retention to decline significantly. Studies using a 25-nm-MLC (Multi-Level Cell)-NAND have found that after five years storage at 55°C only 75% of the original data was correct. A comparatively moderate temperature increase to 85°C reveals retention dropping to below 10%.




Figure 2

The high electrical field during programming and erasure cause electrons to be torn out of the oxides crystal structure, over time creating leakage paths from the storage gate to the substrate. These cracks in the tunnel oxide allow charge to slowly leak. Read errors increase and finally reject the complete block as a "bad block".

„Every write to a page within a block also stresses the adjacent pages.“

## Impact on retention

The impact of oxide layer degeneration on retention is massive: Retention expectancy of SLC and MLC is typically ten years, but at the end of the life-cycle it drops to just one year. For MLC this point is reached after 3.000 P/E cycles and for SLC after 100.000 P/E cycles. Despite its higher cost, the more durable SLC therefore remains the preferred choice for industrial applications.


Charge state and threshold voltage are key considerations. They rule out lower-priced Triple-Level-Cell-NAND-Chips (TLC), popular for consumer applications, for durable long-term storage in industrial applications. TLC requires eight different charge levels to be able to write and read 3 bits per cell. Here, the degenerative effects are much

quicker. For TLC it only takes 500 P/E cycles before the original retention of one year drops to three months.

Understanding that fewer different charge levels make the storage of information on a NAND chip less vulnerable has led to a commercially and technically very interesting compromise. The “pseudo Single Level Cell” (pSLC) process uses the more cost-effective MLC chip (compared to SLC) only for the first “strong” bit per cell. The pSLC mode is significantly faster than the standard process on MLC-Flash memories and increases the possible P/E cycles from 3.000 to 20.000. At the same time, the endurance of data media is increased by 6.7 with only twice the costs per stored bit.

Figure 3

Comparison of Endurance of SLC, pSLC and MLC



## Stressed flash memory

Aging of memory cells is accelerated by erasure. Block erasures are however required in order to write. The danger of a deceptive conclusion is high, that in a pure reading application – for example a boot medium – long-term data security is ensured due to its

extended retention. Unfortunately, this is a clear misconception. Other circumstances can cause read errors and indirectly contribute to wear and tear of NAND cells. During each writing process, stress arises to the cells close to the cell to be programmed, manifesting



itself in a slightly increased voltage (Program Disturb). Any reading causes stress as well (Read Disturb). As the neighbouring pages collect voltage, over time the stored potential in these cells increases. Read errors are generated, which after deletion of the block disappear again. The lower voltage causes the effect to be less strong for reading than for writing, but bit errors occur nevertheless. To a certain limit these can errors be com-

pensated through the error correction code (ECC), but result in a backup copy of the weak block. The effect is particularly strong in applications that repeatedly read the same data. It has to be taken into consideration that even inside a memory medium, which is only used for reading, blocks have to be deleted and pages written regularly as part of the error correction. As a result, a medium that is only reading ages as well.

## The measurement of endurance

*"Only the declaration of the workload allows a comparison of different products. Often this information is missing in data sheets."*

Two metrics are used by manufacturers to measure the lifetime of a flash memory device:

- Terabytes Written (TBW) (Figure4): TBW indicates the total amount of data that can be written during an SSD's lifetime
- Drive Writes Per Day (DWPD): DWPD indicates the amount of data that can be written every single day during the warranty period

These benchmarks are extremely complex, so developers, customers and users have no choice but to rely on the manufacturer's specifications. The big challenge is to determine whether the specifications have any significance for the required application, as the values very much depend on the type of workload applied during the test. Measurements made by Swissbit using their 480-GB-

SSD resulted in an endurance of 1.360 TBW (Sequential Writing), 976 TBW (Client Workload) and 240 TBW (Enterprise Workload) depending on the measuring process. The Client Workload was based on the user behaviour of a PC user and generated mainly sequential access. The Enterprise Workload simulated the behaviour of a server in a multi-user environment, which generated 80 % random access.

Guidelines for endurance testing are set by the standardization organization JEDEC and should ensure the comparability of products and manufacturers. Workload specifications, however, are often not included within datasheets. Caution is advised from manufacturers claiming phenomenal endurance values, as these are often based on sequential writing only used in few applications. Endurance values of flash solutions can easily differ by a factor of 10 for Sequential Writing and Enterprise Workload.

$$TBW = \frac{\text{Drive Capacity (GB)} \times \text{Cycles}}{WAF} \left( \frac{TB}{1000 GB} \right)$$

Figure 4

TeraBytes Written Equation

# Methods of controller and firmware optimization – “Internal affairs”

*"For non-sequential writing, page-based mapping is a better fit."*

Knowing the limiting effects we now take a closer look at what manufacturers do, to better use what the NAND chips have to offer. pSLC has already been mentioned and it leads the way: the key to exceptional performance or durability is the controller, respectively the firmware. It is important to understand that erase, writing and reading are not only triggered by the actual application, but also by numerous processes caused by controllers and firmware without the application or user noticing.

A number of internal mechanisms come into play.

- Error Correction: After bit errors have been corrected, the block where the corruption occurred is copied and then deleted
- Garbage Collection: re-copying for the release of blocks
- Wear Levelling: Applications which predominantly write only on few logical addresses would quickly wear out the physical blocks hosting these logical addresses. The wear levelling introduces a logical to physical mapping which allows that any logical block can be assigned to any physical location and thus the physical array will be evenly utilized and stressed

Mapping between the logical and physical address makes data storage possible in the first place. The described processes complement this mechanism. The ratio between the user data coming from the host and the actual data size written onto the flash memory determines the efficiency of the flash medium’s controller. The value is expressed by the write amplification factor (WAF). (Figure 5)

Reducing the WAF is the key for longer endurance. WAF is also influenced by workload factors, such as the difference between sequential and random access or the size of the data blocks in relation to page and block sizes. Consequently, the firmware also determines the suitability of a flash medium for a specific application.

$$WAF = \frac{\text{Total Bytes Written to Flash}}{\text{Total Bytes Written from Host}}$$

Figure 5

Write Amplification Factor

## Increasing efficiency

An even greater understanding can be attained by delving deeper into the operating principles of flash memory. Pages of a block of cells must be programmed in succession and only complete blocks can be deleted. In the standard process, mapping between the logical and physical address refers to blocks. This is very efficient for sequential data as

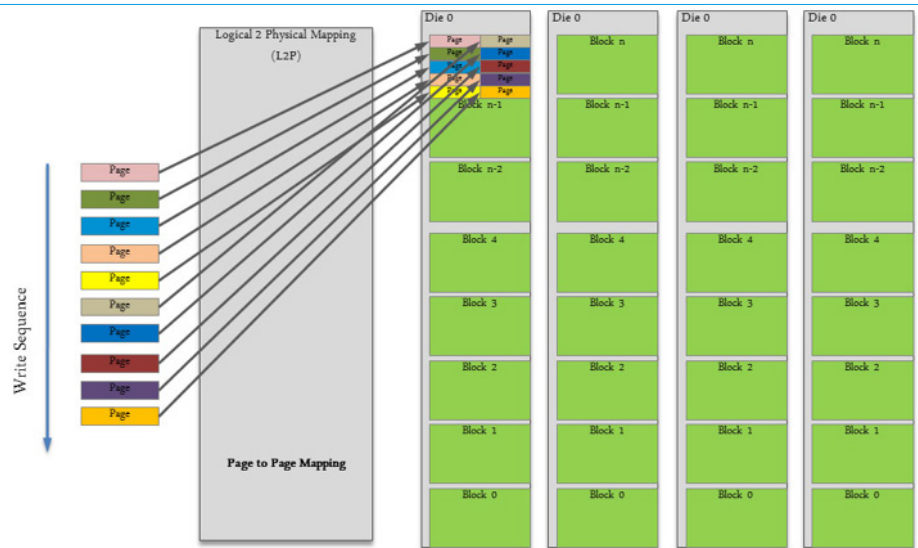
the pages of a block are written in succession. For example, block based mapping is ideal for an application such as continuously collected video data.

For random data, however, pages are written in many different blocks. Here, a complete block has to be deleted per page for each internal re-programming. This results in a high WAF and an endurance-decrease. Page-based mapping is therefore better suited for non-sequential data where

the firmware ensures that data of different origins is saved sequentially on the pages of one block. (Figure 6) The number of deletions is thereby reduced – which has positive consequences for endurance – and the write performance is increased. Whilst page-based mapping increases the allocation table of the FTL (Flash Translation Layer), manufacturers compensate that with integrated DRAM. Therefore page-based mapping can be beneficial.

Figure 6

*Page based mapping:  
Random writes from  
different locations can be  
combined in one flash  
block.*



## Over-provisioning

Further benefits of page-based mapping occur where the degree of utilization of the data medium forces the WAF up. The more data is stored on the flash medium, the more bits have to be moved back and forth by the firmware. Manufacturers prevent problems with overloaded data media simply by over-provisioning, increasing the flash area which is reserved only for internal activities. Conventionally, that is the 7 percent difference between decimal and binary values of a Gigabyte specification.

Increasing the area reserved for management task to 12% creates a surprising effect.

An endurance comparison (TBW for enterprise workload) between two SSDs of identical hardware revealed that the Swissbit model X-60 durabit™ with 240 GB achieved a value almost double that of a 256 GB configuration. A comparison of the impact of the DRAM on endurance showed that the 240 GB durabit™ version was even 10 times higher compared to the standard version with 256 GB. It is important to point out that similarly to using MLC as pSLC, a significantly positive endurance effect can be achieved by sacrificing memory capacity or by applying over-provisioning. (Figure 7)

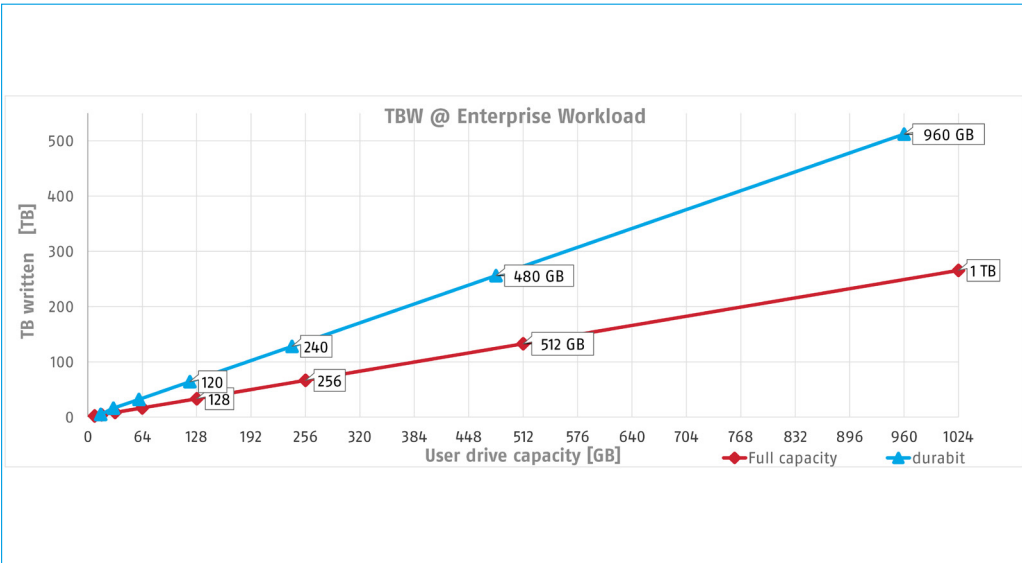


Figure 7

All new MLC based Swissbit SSDs utilize a DRAM supported Page Based Mapping. (X-60, X-60m2, F-60 and others). *durabit™* MLC products use higher overprovisioning for up to 20% increased performance and 100% more endurance.

## Data maintenance

*"Even though the main aging effects are triggered by erase cycles, there is a risk for read-only applications if degrading bits are not refreshed."*

Whilst error correction and wear levelling are mechanisms that are also used in universal flash products, manufacturers of high-quality industrial SDDs or flash memory cards will go even further to prevent data loss and system failures. By combining different mechanisms such as ECC monitoring, read disturb management and auto read refresh, all stored data is monitored and periodically refreshed as required and system failures are prevented in advance. A key point is that data integrity should be ensured without involving the host application, requiring these processes to run autonomously within the memory card and not only in the usual instance of cumulative bit errors following read requests by the host application.

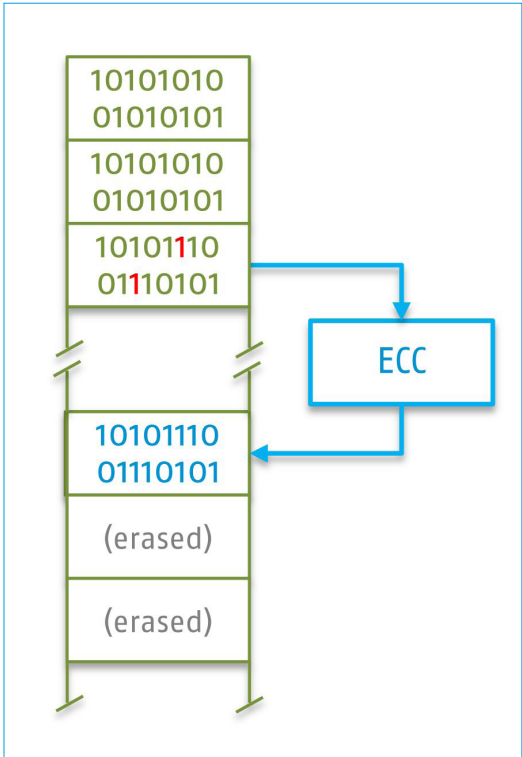


Figure 8

Gradual data loss is prevented by data care management. All written blocks are read in the background and copied, repaired and rewritten in case of too many bit errors.



Advanced data care management searches independently of requests by applications for potential errors (Figure 8). All written pages, including the firmware and the allocation table of the FTL, are read in the background and refreshed as required. The different triggers for this preventative error correction include:

- a defined number of repeated switching on
- the number of P/E cycles
- the read data volume
- read repetitions/re-readings
- increased temperature

## Conclusion

*"Interestingly a slight reduction of usable capacity in exchange of increased overprovisioning, nearly doubles endurance."*

A thorough understanding of the characteristics of flash technology is the key for the selection of the most suitable storage solution for industrial applications. Indeed, other criteria should not be ignored: such as power-fail protection mechanisms, particularly robust processing and the specification

for an extended temperature range.

Long-term availability of modules that were meticulously meticulously qualified for a particular application is an important, yet often overlooked selection criteria. For this reason, 3D-NAND was not discussed in this paper because:

- the technology is still too new
- innovation cycles and design changes are currently too ad-hoc for the life cycles of industrial products
- the chips are all optimized for the TLC consumer sector and not specified for extended temperature ranges
- empirical values for endurance and retention of these chips doesn't exist yet

As this paper has illustrated, the key task for manufacturers of industrial flash products is to optimize these values.

Every single Swissbit product is tested before it leaves the factory in Berlin. High-

est quality standard and long-term availability of devices with a fixed BOM make Swissbit the perfect partner for storage that fits the requirements of industrial embedded systems.

## Annexe

When selecting flash memory for an industrial application, the evaluation process should be preceded by requirements engineering, which also includes the type and

quantity of data to be written and read by the application. The following are some typical questions and answers that will help you choose a suitable storage solution.

### Questions and Answers:

- What particular physical requirements are needed with regard to vibration-resistance and temperature range of the medium? – Industrial flash memories should go through high-quality testing of materials and processing to specify and qualify proven properties.
- Is the memory exposed to high temperatures over longer periods of time? – High temperatures weaken the readability of cells faster. Therefore, it is best to choose a product with data care functions that regularly refreshes data.
- Do you intend to store a lot of data on the data medium over a long period and does it need to be maintained for a long time? – An SLC product is best for this type of application.
- Does the application mainly read? – We recommend a medium with data care management that refreshes data regularly.
- Does the application mainly write? – We recommend a product with page based mapping and DRAM supported FTL.
- Will the full capacity of the memory be used? – Where usage is intense, the controller needs space for internal processes – over-provisioning extends endurance.
- Which workload has been used to specify TBW and DWPD? – Data media can only be compared using the same workload benchmark.
- Is increased data loss protection required? – Data care management and power fail protection should be included for particular critical applications.
- How long will the medium be available for? – Long-term availability should be warranted by the manufacturer to allow replacement without the need for requalification.

Contact

Swissbit AG  
Industriestrasse 4  
9552 Bronschhofen  
Switzerland

[swissbit.com](http://swissbit.com)  
[industrial@swissbit.com](mailto:industrial@swissbit.com)



**Author**  
Ulrich Brandt  
Direktor Marketing  
Swissbit AG

Tel +41 71 91303 48  
E-Mail [ulrich.brandt@swissbit.com](mailto:ulrich.brandt@swissbit.com)

Flash-technologies  
in comparison

© Swissbit AG 2017  
Edition 2/0617

