

swissbit®

X-500 / X-55 Series SLC vs. EM-MLC

White Paper

BU: Flash Products
Date: April 17, 2014
Revision: 1.00
File: WhitePaper_X-5xx_EM-MLC vs
SLC_Rev100.doc



1 Overview

Swissbit's latest SSD generation is available in 2 versions:

1. **X-500: based on SLC NAND flash (Single Level Cell)**
for applications with many write operations, requiring highest reliability
2. **X-55: based on EM-MLC NAND Flash (Endurance Managed Multi Level Cell)**
for cost driven applications with lower write workload and lower data retention and performance requirements

This document compares properties of different flash types and explains possibilities to maximize the endurance.

Typical characteristic values of different 24nm NAND flash types are:

	used in X-500	Used in X-55	<i>not used</i>	
	SLC	EM-MLC	MLC	TLC
Endurance (PE cycles)	100,000	30,000	3,000	<1,000
Data retention initial (years)	10	5	10	<2
needed Error correction (bits)	8 / sector	40 / 2 sectors	40 / 2 sectors	96 / 2 sectors

2 Comparison of SLC, EM-MLC, MLC and TLC NAND flash

The difference of SLC, (EM-)MLC, and TLC NAND flash types is the different numbers of bits that can be stored in a single flash cell (1, 2, or 4 respectively).

A NAND Flash cell is the insulated floating gate of a special FET transistor, containing a number of electrons that varies depending on the voltage applied to the cell. The number of stored electrons can be read out by measuring the conductivity of the transistor's different gate voltages.




Figure 1: NAND Flash cell (FET transistor with floating gate)

In the different flash types the chip must differentiate 2, 4, or 8 different states by their numbers of electrons, respectively.

	SLC (2 states)	(EM-)MLC (4 states)	TLC (8 states)
0		00	000
		01	001
		10	010
		11	011
1		100	100
		101	101
		110	110
		111	111

Table 1: Schematic diagram of flash cell filling states and respective bit information for different flash types.

When writing to an SLC flash cell it is either fully filled with electrons or kept empty, but MLC and TLC flash cells have to differentiate 4 / 8 different stages by their loading and allows for far less tolerance in the write and read operation. The actual loading of the flash cell has a Gaussian like distribution:

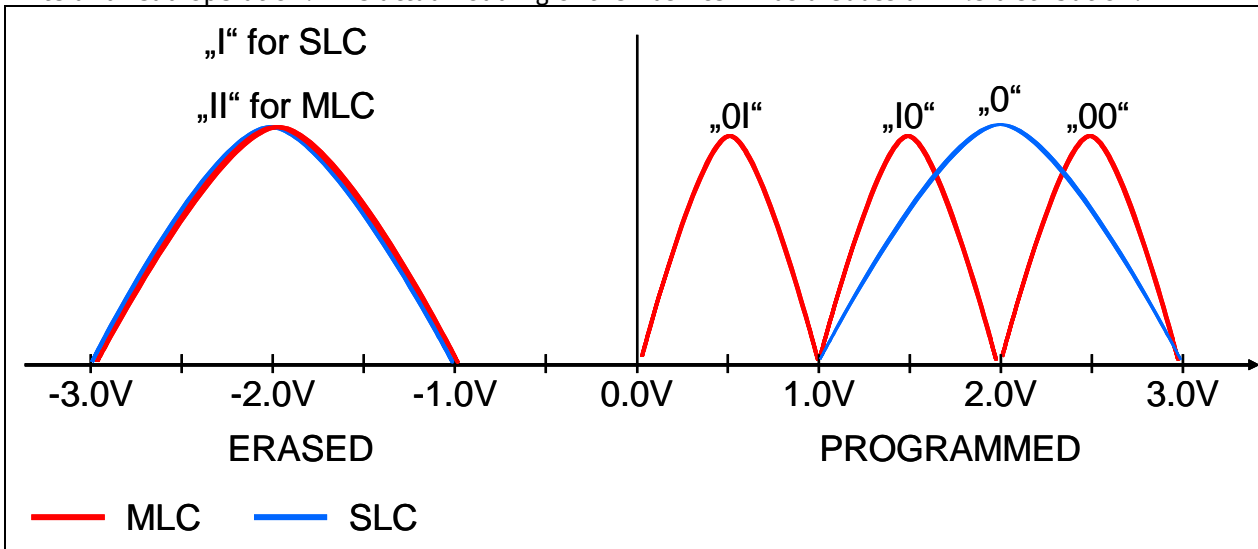


Figure 2: Schematic distribution of states by gate voltage for SLC and MLC Flash

Especially in MLC and TLC flash the states can overlap, so bits can be read out incorrectly (bit errors), and need to be corrected by an ECC (error correction code)

The below effects generally define data integrity of a flash cell:

- 1) **Data retention** specifies the period during which data remains good (i.e. within ECC limit) after being written to the flash. It depends on the number of PE (program erase) cycles and ambient temperature.

- 2) **Program erase cycles:** Every program/erase cycle of a flash cell **damages the insulator** of the floating gate a little, making it more and more porous. With every PE cycle, the data retention reduces until, at the end of the flash cell's life time, the floating gate cannot keep the electrons any more.
- Also, the threshold voltages of the different states change with increasing number of PE cycles and at some point the actual threshold will cross the pre-set threshold of another stage (see below graphs for illustration) and cause read errors
- Typically, the specified number of PE cycles are 100,000 for SLC flash but only 3,000 for MLC flash.




Figure 3: Shift of the read threshold gate voltage for SLC and MLC flash in dependence on the number of PE cycles.

- 3) **Temperature:** With increasing temperature, more electrons escape from the flash cells. In other words, a higher temperature reduces the data retention. The acceleration factor can be described with the Arrhenius equation, with a typical Activation Energy $E_a=1.1\text{eV}$. Normally the data retention is specified at 40°C. Based on these assumptions we get the following acceleration factors and data retention times depending on temperature:

Temperature	40°C	55°C	70°C	85°C	100°C
Acceleration Factor	1	6.4	35	168	706
Specified retention	10 years	565 days	103 days	11 days	5 days

Table 2: Temperature dependence of the data retention, specified for 10 years at 40°C

In conclusion, if a flash storage device is stored at high temperature data retention is reduced drastically.

- 4) **Read Disturb:** In order to read a flash page, a small voltage has to be applied on the previous pages in the block which, in the case of very frequent access to the same page can lead to unintended soft-programming of surrounding pages.

2.1 EM-MLC compared to standard MLC flash

EM-MLC flash is optimized for a larger number of program-erase cycles.

It is the same flash cells as standard MLC flash but is written/read with different parameters. In EM-MLC flash fewer electrons are injected into the flash cells, so the insulator of the flash cells is damaged less in each program/erase operation compared to conventional MLC. Due to the lower number of electrons, the initial data retention for EM-MLC is lower but decreases far more slowly (MLC: decrease of data retention after 300 PE cycles, EM-MLC: after 3,000 PE cycles). The total number of specified PE cycles of EM-MLC is 30,000 while the equivalent MLC flash has only 3,000 PE cycles.

The X-55 SSD family uses EM-MLC Flash for maximum endurance compared to other MLC based SSDs. The below table provides an overview on the data retention (in years) by Flash type in relation to the number of PE cycles.

PE cycles	0	300	1,000	3,000	10,000	30,000	100,000
EM-MLC 24nm	5	5	5	5	1	0.25	-
standard MLC 24nm	10	10	3	1	-	-	-
SLC 24nm	10	10	10	10	10	4	1
MLC 19nm	5	3	1	-	-	-	-

Table 3: Typical data retention (years) at 40°C

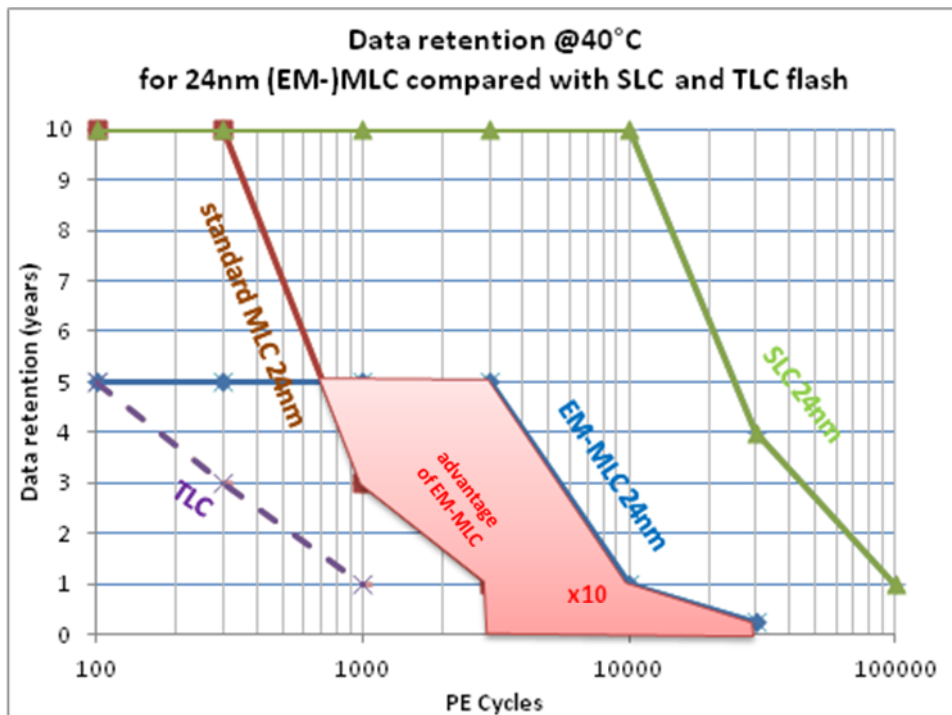


Figure 4: Typical data retention at 40°C

The **initial data retention** of EM-MLC flash is only **half of the standard MLC value**, but the **number of program/erases cycles is 10 times larger**.

After 700 PE cycles the data retention of EM-MLC is longer than for standard MLC Flash.

Nevertheless the retention behavior and maximum number of PE cycles of SLC flash is better in any case.

3 Methods to increase data retention and endurance

The X-500 / X-55 series use different mechanisms in order to increase endurance or data retention:

- 1) **Wear leveling:** The controller counts the number of erase operations for each block. If the erase counter of one block is much larger than another block, both blocks are exchanged. This guarantees a nearly even wear of all flash cells in the SSD.
- 2) **Error correction code:** The X-55's controller and firmware can correct up to 40 bit errors in 2 adjacent sectors
- 3) **Read retry:** If the data cannot be corrected, the controller re-reads the data with different read thresholds which can allow to retrieve data already corrupted through read disturb or data retention related effects.
- 4) **ECC Monitoring:** The controller monitors the error bit levels in each read operation; when it reaches the preset threshold value, it automatically performs a data refresh in order to "restore" the correct charge levels in the cell. Through the refresh, the data is reset to its original, error-free state, and the data retention is increased.
- 5) **Early Retirement:** The Early Retirement mechanism flags blocks as bad blocks just before they are physically broken and hence prevents data corruption. This works by constantly analyzing the bit error rate of the data that is read out and comparing them against a pre-set threshold. Once that is reached, the above actions are taken.
- 6) **Data Scrambling:** The controller stores the data scrambled to the flash, i.e. combined with pseudo random data. This guarantees that the number of "0"s and "1"s is similar in each page and the write current is leveled over the whole page, regardless of the data written by the host.
- 7) **Overprovisioning:** Swissbit's X-55 reserve 6.7% of the physically available capacity for overprovisioning. This reduces the total capacity available to the user but significantly increases the sustained random write speed and endurance.
The 6.7% of the capacity is used as a kind of additional cache that the SSD can use for garbage collection which helps reducing the write amplification factor (i.e. total data written to the flash divided by total data received from host) which increases the SSDs life time.
- 8) **TRIM:** The TRIM command is supported in more recent operating systems (from Windows 7). It allows the OS to inform the SSD which data is not in use in the file system any more so the SSD does not need to consider this data in the wear leveling process. This significantly increases performance and life time of the SSD. It is important to ensure that both operating system and SSD support TRIM.

4 S.M.A.R.T. Information

The S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) Information can be read out using the free SwissBit Life Time Monitoring (SBLTM) tool which also provides additional information such as max/min/current system temperature, average erase count, erase life time status or spare block status. The different maximum number of program erase cycles of SLC and EM-MLC flash is of course considered by the tool when predicting the SSD's life stage.

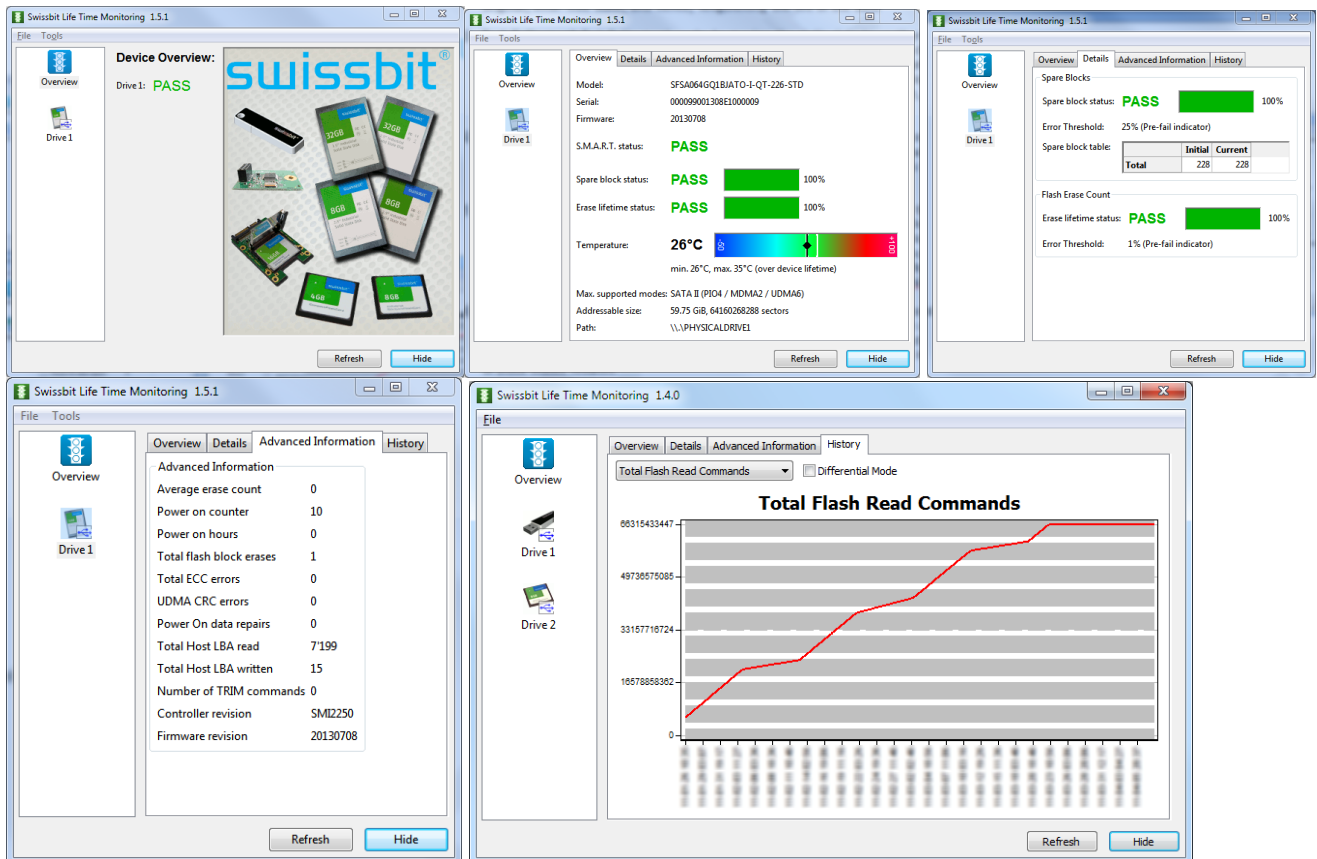


Figure 5: Print screens from the SBLTM tool. They show the interpretation, values, and history of the S.M.A.R.T. information

Please ask your Swissbit representative for this free tool or for an API to implement it into your own application.

5 Document History

Table 4: Document Revision History

Date	Revision	Details
01-October-2013	0.10	First DRAFT
12-February-2014	0.20	
17-April-2014	1.00	First release revision

Disclaimer:

No part of this document may be copied or reproduced in any form or by any means, or transferred to any third party, without the prior written consent of an authorized representative of Swissbit AG ("SWISSBIT"). The information in this document is subject to change without notice. SWISSBIT assumes no responsibility for any errors or omissions that may appear in this document, and disclaims responsibility for any consequences resulting from the use of the information set forth herein. SWISSBIT makes no commitments to update or to keep current information contained in this document. The products listed in this document are not suitable for use in applications such as, but not limited to, aircraft control systems, aerospace equipment, submarine cables, nuclear reactor control systems and life support systems. Moreover, SWISSBIT does not recommend or approve the use of any of its products in life support devices or systems or in any application where failure could result in injury or death. If a customer wishes to use SWISSBIT products in applications not intended by SWISSBIT, said customer must contact an authorized SWISSBIT representative to determine SWISSBIT willingness to support a given application. The information set forth in this document does not convey any license under the copyrights, patent rights, trademarks or other intellectual property rights claimed and owned by SWISSBIT. The information set forth in this document is considered to be "Proprietary" and "Confidential" property owned by SWISSBIT.

ALL PRODUCTS SOLD BY SWISSBIT ARE COVERED BY THE PROVISIONS APPEARING IN SWISSBIT'S TERMS AND CONDITIONS OF SALE ONLY, INCLUDING THE LIMITATIONS OF LIABILITY, WARRANTY AND INFRINGEMENT PROVISIONS. SWISSBIT MAKES NO WARRANTIES OF ANY KIND, EXPRESS, STATUTORY, IMPLIED OR OTHERWISE, REGARDING INFORMATION SET FORTH HEREIN OR REGARDING THE FREEDOM OF THE DESCRIBED PRODUCTS FROM INTELLECTUAL PROPERTY INFRINGEMENT, AND EXPRESSLY DISCLAIMS ANY SUCH WARRANTIES INCLUDING WITHOUT LIMITATION ANY EXPRESS, STATUTORY OR IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

©2014 SWISSBIT AG All rights reserved.