



CTAN006: Endurance Models for Cactus Technologies Industrial-Grade Flash Storage Products

Covered Products: All

1. Introduction

Since endurance for NAND flash-based storage devices are limited by the number of erase cycles than can be performed on each NAND block before it becomes unreliable for data storage, it is important for designers to know and understand the concepts and mathematical models for estimating the expected write endurance for Cactus Technologies industrial-grade flash storage products.

2. The Problem

The design of the NAND flash memory has the limitation of a finite erase cycle for each NAND block, therefore write/erase endurance for the device is fundamentally limited by the cycle limit for each NAND block. The controller may mark a block as unusable when:

1. The controller detects an uncorrectable read or write error when reading from or writing to the block.
2. The erase count of a particular NAND block has reached a predefined threshold.

Current SLC NAND technologies as used in all Cactus Technologies industrial-grade flash storage solutions enable each cell to be erased for 100,000 cycles.

NAND blocks with very high erase counts are often unreliable during write/erase. In many cases for NAND-based flash storage devices, the usable life is defined as the time before the host can no longer reliably write data to the device.

3. The Solution

Advanced error detection and correction algorithms can extend device endurance by correcting multiple read errors when they occur due to device aging. The on-board intelligent controller on Cactus Technologies industrial-grade flash storage products implement a powerful ECC to correct multiple errors per sector.

To increase the number of effective write/erase cycles, many on-board controllers reserve a percentage of total available NAND blocks as "reserved blocks". When the controller detected an unrecoverable write/erase error, a block from the reserved block pool is used to replace the failed block. This operation is similar to sector remapping schemes implemented by modern magnetic disk drives, and is completely transparent to the host. Defect block remapping increases data reliability and prolong usable life. All Cactus Technologies industrial-grade flash storage devices implement defect block remapping with reserved NAND blocks on the device.

In many applications, data writes (translated into page programming cycles) occur less often than data reads (translated into page read cycles). Therefore there are data blocks on the devices which are more likely to be erased than others due to writes/updates from the host. If the average erase count across the entire chip or the entire device can be exploited by swapping blocks containing static data (i.e. data not often updated) with blocks containing data often updated, then overall endurance can be significantly increased. This procedure is termed "wear-leveling". It is implemented in the on-board intelligent controller for all Cactus Technologies industrial-grade flash storage products.

Cactus Technologies industrial-grade flash storage products specifies endurance at 2 million logical erase cycles. This rating is derived from typical usage, and may differ among applications as described in the next section.

4. Endurance Estimations

The following scenarios show that the expected endurance changes primarily in terms of usage pattern. The device capacity also plays a part in endurance.

The following estimation are based on the following assumptions:

1. Wear leveling forces write/erase be always performed to a different physical NAND flash block each time and moves data across the whole device.
2. Read operations does not affect device endurance and is ignored.
3. ECC errors are completely correctable, therefore block replacement is not required.

A general model for endurance estimation may be derived as follows:

D = data size (kilobytes)

b = block size ((kilobytes)

T = user accessible blocks per NAND flash IC

n = number of NAND flash IC present on the device

e = endurance for a single NAND block

The number of blocks erased **B** in terms of data size during write and flash block size is evaluated with the following expression:

$$B = \frac{D}{b}, B \geq 1$$

The erase behavior of full block erasure means that if **D** < **b**, then **B** = 1, therefore **B** ≥ 1 for all calculations.

Since there is **nT** number of blocks on the device depending on the configuration, the number of erase cycles **t** required to completely erase every NAND block on the device once is evaluated with the following expression:

$$t = \frac{nT}{B} = \frac{nTb}{D}$$

The device may be “retired” when every block reaches the endurance count **e**. Device endurance **E** is evaluated with the following expression:

$$E = e^{\left(\frac{nTb}{D}\right)}$$

The model above shows device endurance is related to device capacity and data size per write operation.

The following scenarios with two hypothetical industrial-grade flash storage devices illustrate how the endurance model helps in estimating device endurance.

	Device A	Device B
Capacity	2GByte	64MByte
Number of NAND flash IC (n)	2	2
User accessible blocks per IC (T)	8032	2013
Block Size	128Kbyte	16Kbyte
Erase cycle endurance	100,000	100,000

Scenario 1: Sequential data writes to Device A

starting from logical sector 0 until the device is filled then repeated. Each write operation writes 128KByte (equal to NAND flash block size) to the device.

$$\begin{aligned}
 E &= e^{\left(\frac{nTb}{D}\right)} \\
 &= 100000 \times \left(\frac{2 \times 8032 \times 128}{128}\right) \\
 &= 1606400000
 \end{aligned}$$

The model shows that if small data blocks are written to a large capacity device, device endurance **E** is determined by the number of user blocks.

Scenario 2: Random data writes to Device A.

The data size is equal to 5% of card capacity every write session. The remaining capacity holds static data. Each write operation writes 100MBytes or 80 NAND blocks.

$$E = e^{\left(\frac{nTb}{D}\right)}$$

$$= 100000 \times \left(\frac{2 \times 8032 \times 128}{(80 \times 128)}\right)$$

$$= 20080000$$

This scenario shows that for typical applications, it is possible for the device to exceed the rated 2 million logical erase cycles.

Scenario 3: Random data writes to Device B.

The data size is equal to 5% of card capacity every write session. The remaining capacity holds static data. Each write operation writes 3.2MBytes or 200 NAND blocks to the device.

$$E = e^{\left(\frac{nTb}{D}\right)}$$

$$= 100000 \times \left(\frac{2 \times 2013 \times 16}{(200 \times 16)}\right)$$

$$= 2013000$$

This shows even for devices built with small block NAND IC as in Device B, a 2 million logical erase cycle rating is achievable.

Scenario 4: Overwrite all blocks in Device A. The data size is equal to the the device capacity.

$$E = e^{\left(\frac{nTb}{D}\right)}$$

$$= 100000 \times \left(\frac{2 \times 8032 \times 128}{(2 \times 8032 \times 128)}\right)$$

$$= 100000$$

This scenario shows that wear-leveling will no longer be effective if the whole device is overwritten as it forces the on-board intelligent controller to

erase all user blocks every time. The total endurance may be extended by the on-board ECC.

A summary table of endurance vs. usage pattern may be drawn as below:

Scenario	Write Intensity	Endurance
1	Low	~ 1.6 Billion
2	Typical	~ 20 Million
3	Typical	~ 2 Million
4	High	~ 100,000

5. Conclusions

From the mathematical model in Section 4, device endurance depends largely on device capacity and usage patterns, and may change due to different usage patterns.

However, it must be emphasized that due to manufacturing variations found in NAND devices, the mathematical model above should only be used as a reference. Designers should always implement data integrity checking and correction systems in their designs with Cactus Technologies industrial-grade flash storage products to achieve the required level of data integrity and system reliability.

6. Version History

Version	Date	Change
1.00	December 12, 2006	Initial Version
1.01	May 25, 2007	Specified "SLC NAND" flash for all Cactus Technologies products
1.02	June 3, 2008	Minor edits