



## **White Paper**

# **M-Systems' Fast Flash Disk LIFE EXPECTANCY CALCULATION BASED ON FLASH ENDURANCE**

Written by: Nimrod Wiesz  
W.W. Tech Support Manager

**APRIL 02**  
45-SR-003-05-8L

## 1 Introduction

There are several ways to calculate the life expectancy of a solid-state flash disk like M-Systems' Fast Flash Disk (FFD). One way is to analyze the reliability factors of the individual components and then calculate an expected Mean Time Between Failure (MTBF). The MTBF calculations are addressed in a different document per each FFD model and specific capacity, which can be supplied upon request. This white paper will present the life expectancy of the FFD based on the endurance of the flash memory inside the FFD.

## 2 Background

Many customers are worried about using a flash disk in write-intensive applications because they have heard that flash has a limited life expectancy. It is true that flash has some limitations, however the life expectancy of flash is one of the most misunderstood parameters of this non-volatile storage technology. All types of flash memory have a specified life expectancy. It is not measured in years. Instead the life expectancy is measured in the number of erase cycles that each erase unit in the flash device can undergo. Therefore we need to find a way to correlate the number of erase cycles with the life expectancy of the FFD measured in years.

It is important to understand that if the flash is not managed properly then this parameter can severely affect the overall life of the flash disk. M-Systems prides itself in having the best, patented, technology for ensuring data reliability and extending the life of the flash. This is the TrueFFS<sup>®</sup> technology.

**Note:** *M-Systems has manufactured and shipped thousands of FFDs for the past seven years. So far no FFD was returned because the flash was worn out since FFD was first introduced to the market in 1997.*

## 3 M-Systems' Technology

There are several methods for managing flash memory and using it to emulate a disk drive. Simple algorithms map a logical sector to a fixed physical location. It is easy to show that such an algorithm will quickly cause the flash to wear out when an application updates the same sectors over and over again<sup>1</sup>. Therefore a simple flash management algorithm could suffer catastrophic failures after only several hundreds of thousands of file operations. The life of the device would depend on the number of erase cycles guaranteed by the flash vendor and the frequency at which each specific group of sectors was updated.

On the other hand, M-Systems' TrueFFS<sup>®</sup> includes several algorithms to ensure that writes and erases are evenly spread across the entire flash array no matter which logical sectors are written. These algorithms implement something called wear leveling<sup>2</sup>. All flash vendors recommend using some form of wear leveling in order to ensure the maximum life of the flash. Wear leveling will also

---

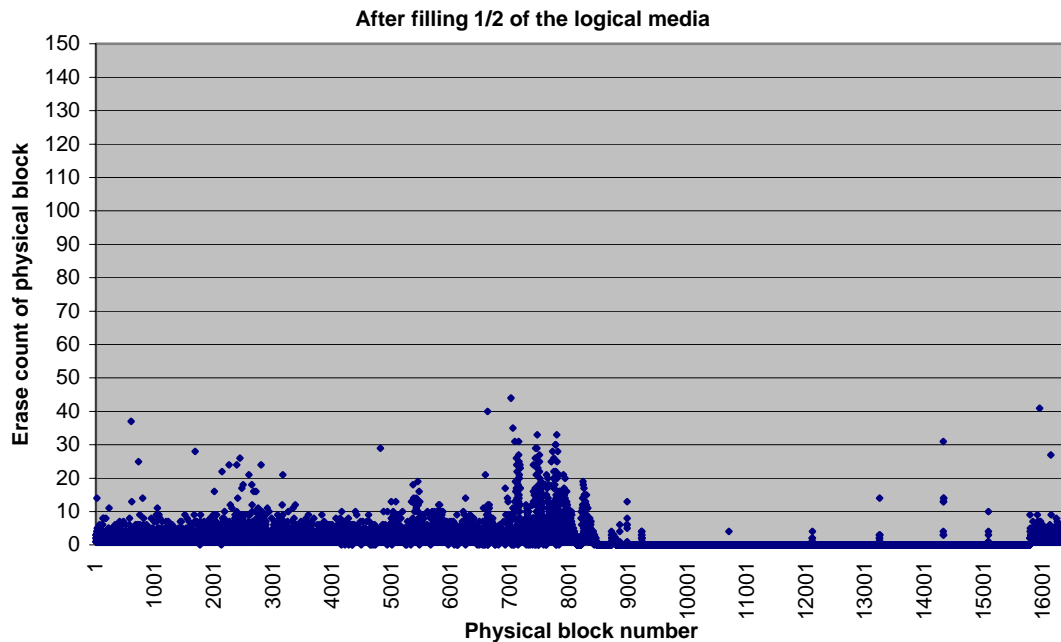
<sup>1</sup> Updating the same group of sectors is a very common scenario. All file systems need to maintain some data that describes the allocation of sectors to files. This data is located in a certain area of a disk drive. For example a FAT file system will update the FAT every time a file is extended or concatenated. The FAT resides in sequential sectors located at the beginning of the media.

<sup>2</sup> A complete description of the algorithm is beyond the scope of this document.

delay the onset of certain failure mechanisms in the flash<sup>3</sup>. These failure mechanisms can cause entire erase units to become inoperable. When wear leveling is used, the erase cycle limit of the flash is increased beyond the minimum specified by the flash vendors.

### 3.1 Wear Leveling Example

In order to show how wear leveling works, an FFD was tested with a special version of firmware<sup>4</sup>. This firmware reported the physical address of the flash erase unit that was erased. Figure 1 Shows the erase count after the media was low level formatted and half of the virtual (i.e. emulated) disk capacity was filled. This figure also shows that when the media is completely erased, the physical erase units are allocated sequentially. Figure 2 shows the state of the media when a test program writes data to the other half of the virtual media. Figure 3 shows the state of media after continuing to run the same test as in Figure 2.



*Figure 1 - Media status after filling half media*

<sup>3</sup> Refer to: Cappelletti, Paulo (June 1999). *Flash Memories*. Kluwer Academic Press, ISBN 0792384873.

<sup>4</sup> This firmware is not available for testing by customers because it requires a specially modified FFD.

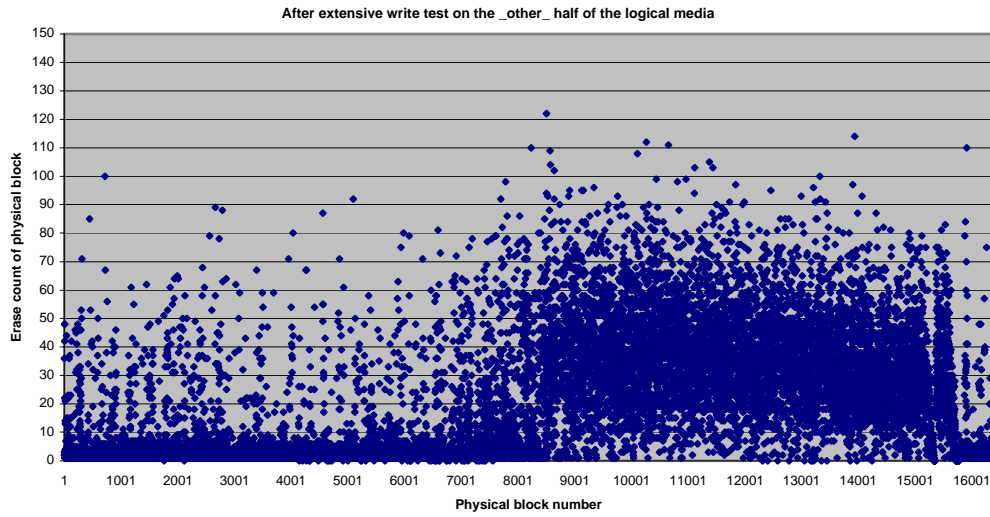


Figure 2 - Media status after continuously updating second logical half

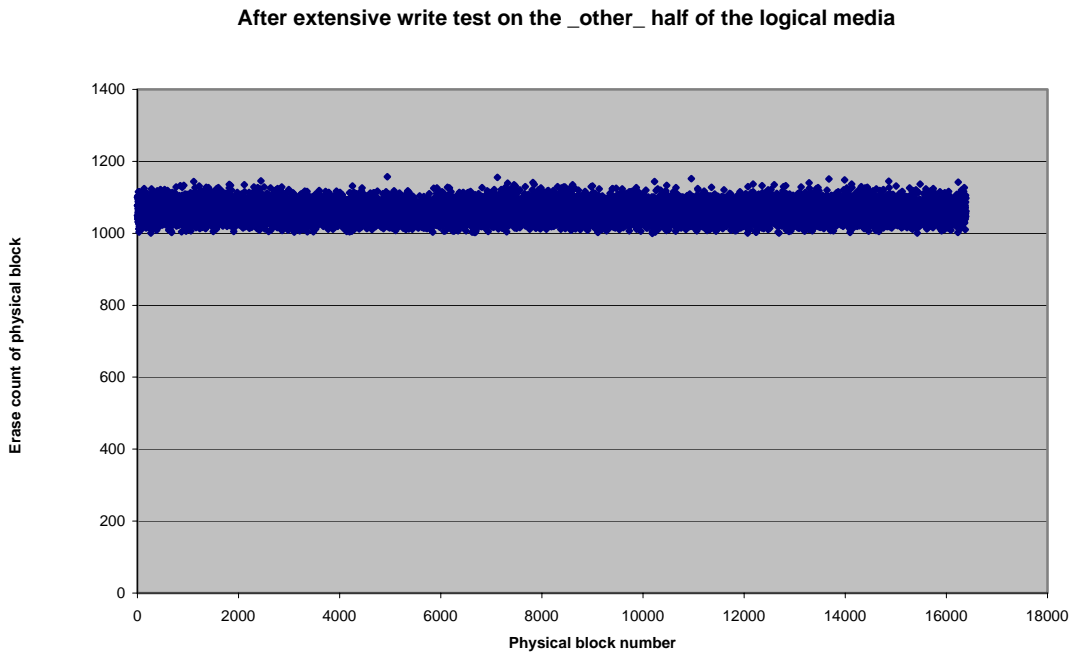


Figure 3 - Media status after continuously updating second logical half (continue)

These figures show that the algorithm does the following:

1. Evenly spreads the number of erase cycle across all the erase units in the flash memory.

2. Reuses areas of flash that contain static data thereby guaranteeing that the entire media is used.

## 4 Life Expectancy Calculation

### 4.1 M-Systems' FFD Life Expectancy Calculation

The FFD incorporates M-Systems' TrueFFS<sup>®</sup> technology. Therefore the worst case FFD life expectancy depends on the following:

1. Media size - The larger the media the more flash components in the FFD.
2. Flash endurance - The flash erase cycle limit guaranteed by the flash vendor<sup>5</sup>.
3. TrueFFS<sup>®</sup> overhead factor - This factor represents the wear-leveling overhead.
4. Amount of data written per time unit - The rate at which data is written to the FFD.

$$LifeSpan(time) = \frac{MediaSize \times FlashEndurance \times 0.995}{\frac{AmountOfdataWritten}{UnitOfTime}}$$

### 4.2 Simple Algorithm Life Expectancy Calculation

Compared to the TrueFFS<sup>®</sup> algorithm used in the FFD, many othersolid-state disks implement a completely different method for managing the flash. In these devices logical sectors are mapped to the same physical erase units until these units reach their erase cycle limit. When the erase units reach their limit they are replaced by erase units located in a pool of spare units. Devices that use this type of flash management algorithm require a different equation for calculating the life expectancy of the disk. For these disks there is a strong dependency between the actual sectors written and the life of the device. The life expectancy equation must take into account where the data is located and the size of the replacement pool. The equation presented for the FFD does not hold for these devices. Instead we must use an equation that takes into account the worst-case scenario, i.e. writing the same group of sectors over and over. This scenario is actually the common scenario for disks used with most types of file systems.

$$LifeSpan(time) = \frac{(1 + SizeOfSparePool) \times FlashEndurance}{\frac{AmountOfDataWritten}{UnitOfTime}}$$

$$LifeSpan(time) = \frac{(1 + SizeOfSparePool)}{AmountOfDataWritten} \times FlashEndurance \times UnitOfTime$$

<sup>5</sup>This will give you the worst-case life expectancy. This number can be replaced by the typical erase cycle limit because TrueFFS is used to manage the data on the media. The typical value is larger than the minimum value by a factor of ten.

### 4.3 Examples

1. Assume you are using a 512MB FFD with Toshiba flash components, which are specified for 250,000 erase/write cycles. Also assume that the application is writing enough data to fill the entire disk once a day. This data can be written to the same group of sectors over and over again or it can be written into the entire media. The location does not matter. The only thing that matters is the total amount of data transferred into the FFD. Putting these numbers into the equation will give:

$$\text{Life Span} = \frac{512\text{MB} \times 250,000 \times 0.995}{\frac{512\text{MB}}{1\text{day}}} = 250,000 \times 0.995 \times 1\text{day} = 248,750 \text{ days or } 681.5 \text{ years.}$$

2. Assume you are using a 512MB solid-state disk managed by a simple algorithm. Also assume that it has 1% units in its spare pool. The disk is installed in application that updates a file every second. The update consists of writing 40Kbytes data. The flash components endurance is specified for 250,000 erase write cycles.

$$\text{SizeOfSparePool} = 512\text{MB} * 0.01.$$

$$\text{AmountOfDataWritten (time)} = 40\text{KB/second} = 3,375\text{MB/day}$$

In this case the life expectancy will be:

$$\text{LifeSpan(days)} = \frac{1 + 512\text{MB} \times 1\%}{3,375\text{MB}} \times 250,000 \times 1\text{day} = 453\text{days}$$

Life expectancy = 453 days or **only** one year! Replacing the simple algorithm with TrueFFS® will give a life expectancy of:

$$\text{LifeSpan(days)} = \frac{512\text{MB} * 250,000 * 0.995}{3,375}$$

Life expectancy = 37,736 days or 103 years.

## 5 Conclusion

The life expectancy of the FFD is not limited by the erase cycle limit of the flash nor does it depend on the application's write patterns. The FFD will probably never reach the erase cycle limit of the flash.

## 6 How to Contact Us

**Website:** <http://www.m-sys.com>

**General Information:** [info@m-sys.com](mailto:info@m-sys.com)

**Technical Information:** [Techsupport@m-sys.com](mailto:Techsupport@m-sys.com)

### USA

M-Systems Inc.  
8371 Central Ave, Suite A  
Newark CA 94560  
Phone: +1-510-494-2090  
Fax: +1-510-494-5545

### Taiwan

M-Systems Asia Ltd.  
Room B, 13 F, No. 133 Sec. 3  
Min Sheng East Road  
Taipei, Taiwan R.O.C.  
Phone: +886-2-8770-6226  
Fax: +886-2-8770-6295

### Japan

M-Systems Japan Inc.  
Arakyu Bldg., 5F  
2-19-2 Nishi-Gotanda Shinagawa-ku  
Tokyo 141-0031  
Phone: +81-3-5437-5739  
Fax: +81-3-5437-5759

### China

M-Systems China Ltd.  
25A International Business Commercial Bldg.  
Nanhu Rd., Lou Hu District  
Shenzhen, China 518001  
Phone: +86-755-519-4732  
Fax: +86-755-519-4729

### Europe and Israel

M-Systems Ltd.  
7 Atir Yeda St.  
Kfar Saba 44425, Israel  
Phone: +972-9-764-5000  
Fax: +972-3-548-8666

© 2002 M-Systems Flash Disk Pioneers, Ltd. All rights reserved.

This document is for information use only and is subject to change without prior notice. M-Systems Flash Disk Pioneers Ltd. assumes no responsibility for any errors that may appear in this document. No part of this document may be reproduced, transmitted, transcribed, stored in a retrievable manner or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without prior written consent of M-Systems.

M-Systems products are not warranted to operate without failure. Accordingly, in any use of the Product in life support systems or other applications where failure could cause injury or loss of life, the Product should only be incorporated in systems designed with appropriate and sufficient redundancy or backup features.

Contact your local M-Systems sales office or distributor, or visit our website at [www.m-sys.com](http://www.m-sys.com) to obtain the latest specifications before placing your order.

DiskOnChip®, DiskOnChip Millennium® and TrueFFS® are registered trademarks of M-Systems. DiskOnKey™ and FFD™ are trademarks of M-Systems. Other product names mentioned in this document may be trademarks or registered trademarks of their respective owners and are hereby acknowledged.