M-Systems
Flash Disk Pioneers

Increasing the Life Expectancy of Flash Disks

White Paper

Nimrod Wiesz

JANUARY 2005
45-SR-003-05-8L, Rev 2.0

TABLE OF CONTENTS

Introduction **3**
Applicability **3**
Background **3**
Flash Management **3**
Mapping a Logical Location to a Fixed Physical Location 4
Mapping a Logical Location to a Dynamic Physical Location 4
Life Expectancy Calculation **6**
Flash Disks Using “Erase before Write” Algorithm 6
M-Systems Products Using TrueFFS Technology 7
Examples 7
“Erase before Write” Algorithm 8
TrueFFS Technology 8
Conclusion **8**
How to Contact Us **9**

M-Systems
Flash Disk Pioneers

INTRODUCTION

This white paper evaluates the life expectancy of flash disks based on the endurance of the embedded flash memory. This method can be applied together with alternate methods, such as analyzing the reliability factors of the individual components within the flash disks and then calculating the flash disk expected Mean Time Between Failure (MTBF) rate. In addition, actual MTBF figures can be calculated based on the return of faulty units using Return Material Authorization (RMA) data. Using all these methods provides a true picture of flash disk life expectancy.

APPLICABILITY

This document is applicable to M-Systems products that implements M-Systems flash management technology, including:

- FFD products
- uDiskOnChip products
- DiskOnChip products

BACKGROUND

Many customers are concerned about using a flash disk in write-intensive applications because they have heard that flash has a limited number of write and erase cycles. This limitation is one of the most misunderstood parameters of flash, which is a highly efficient non-volatile memory technology.

All types of flash memory have a life expectancy. It is not measured in years. Instead, it is measured as the number of erase cycles that each erase unit in the flash device can undergo. This document will correlated the write and erase cycles number with the life expectancy, measured in years, of the product in which it is embedded.

One of the key ways to extend flash life expectancy is by proper flash management. If flash is not managed properly, its life expectancy can severely diminish the overall life of the flash disk. On the other hand, using a good flash management tool can substantially increase flash life expectancy to the point where the limit is of no consequence in even the most write-intensive applications. Properly managed flash, for example, is required in surveillance systems on public buses to prevent crime, and in black boxes to continuously record procedures and events. These systems record data non-stop on flash disks, overwriting old recorded data with new data and keeping only the last few hours of recorded data.

FLASH MANAGEMENT

There are several methods for managing flash memory when using it to emulate a disk drive. Two of these methods are discussed below.

Mapping a Logical Location to a Fixed Physical Location

Simple algorithms map a logical sector to a fixed physical location. This method quickly causes the flash to wear out when an application updates the same sectors over and over again. 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 specified area of a disk drive. For example, a File Allocation Table (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.

A simple flash management algorithm could suffer catastrophic failures after only several 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 with which each specific group of sectors was updated.

Mapping a Logical Location to a Dynamic Physical Location

Using an algorithm to dynamically map a logical sector to a physical location can ensures that all write and erase cycles are evenly spread across the entire flash array. This method, called wear-leveling, extends flash life expectancy. Wear-leveling is dynamically applied to all the flash blocks to ensure the greatest effect on increasing flash life expectancy.

More than extending flash life expectancy, wear-leveling also delays the onset of certain failure mechanisms in the flash[1]. 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 flash vendors.

M-Systems' TrueFFS® (True Flash File System) flash management tool, embedded inside its Fast Flash Disk (FFD™) , uDiskOnChip and DiskOnChip families, incorporates one of the most effective wear-leveling mechanisms.

To demonstrate how M-Systems' wear-leveling algorithm works and its results, an FFD[2] was tested with a firmware version designed especially for this purpose. This firmware reported the physical address of the erased flash erase unit. *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 second half of the virtual media. *Figure 3* shows the state of the media after continuing to run the same test as in *Figure 2*. These figures demonstrate that the TrueFFS wear-leveling achieves the target of evenly spreads the number of erase cycle across all the erase units in the flash memory.

[1] For details, refer to *Flash Memories,* Cappelleti, Paulo (June 1999). Kluwer Academic Press, ISBN 0792384873.

[2] Being the fastest M-Systems disk, the FFD was selected to demonstrate M-Systems wear-leveling.

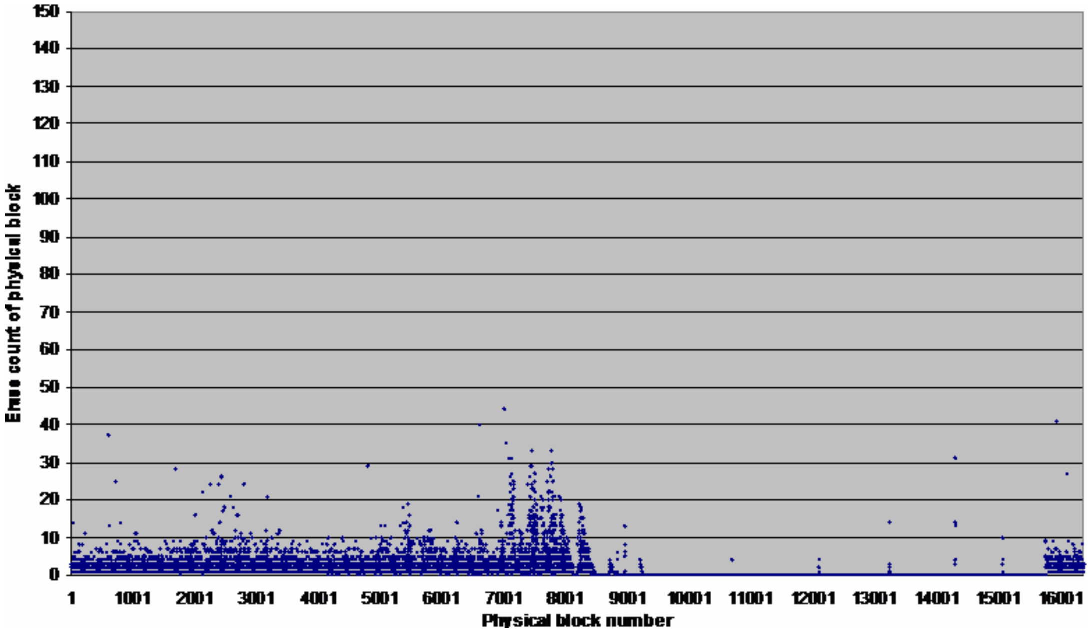

Figure 1: Media Status after Filling Half the Media

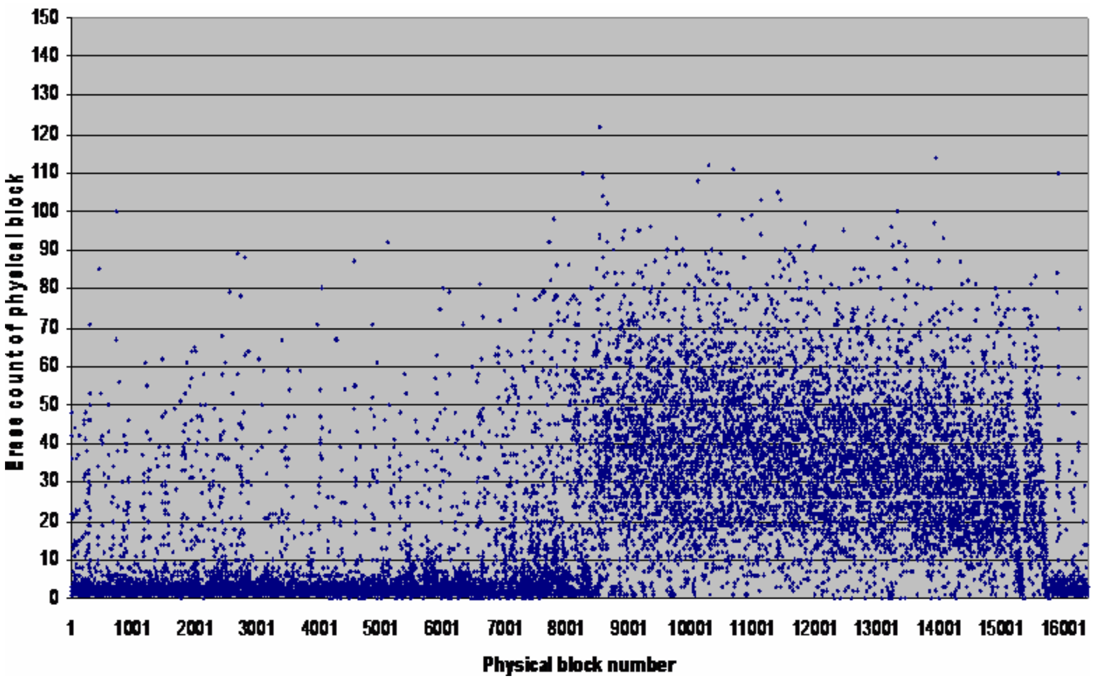

Figure 2: Media Status after Continuously Updating Second Logical Half (Part 1)

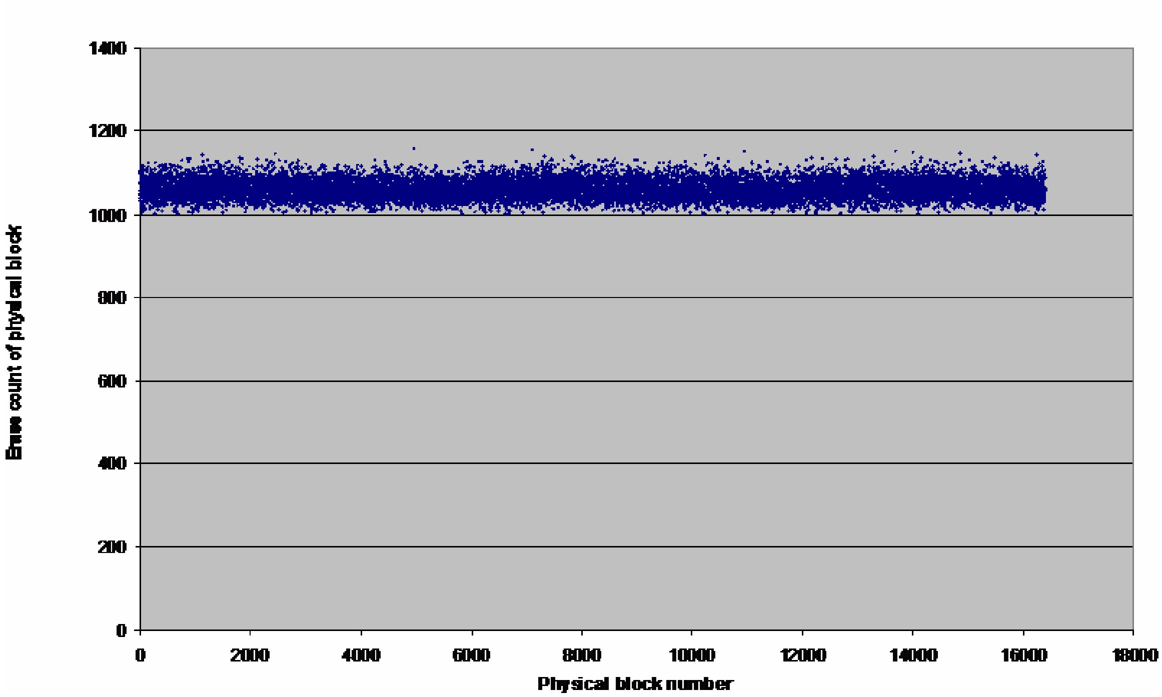

Figure 3 - Media Status after Continuously Updating Second Logical Half (Part 2)

LIFE EXPECTANCY CALCULATION

A number of methods can be used to calculate the life expectancy of the product in which the flash is embedded, depending on the type of flash management used.

Flash Disks Using "Erase before Write" Algorithm

The most common flash management algorithm used is called "Erase before Write". Every time a sector (512Bytes) is written, an erasable block of must be erased. The number of sectors in a block depends on the flash density. In 8 Gbit (8192 MB) NAND flash components, one erasable block contains 256 sectors (128 Kbytes).

When using this method, flash disk endurance is calculated based on the following parameters:

- Number of sectors within erasable block
- Media size (disk capacity)
- Flash disk endurance (write/erase cycles guaranteed by flash disk vendor)
- Amount of data written per time unit (rate at which data is written).

The calculation is as follows:

$$LifeSpan(time) = \frac{MediaSize \times FlashDiskEndurance}{\frac{AmountOfDataWritten}{UnitOfTime} \times No.SectorsWithinErasableBlock}$$

M-Systems Products Using TrueFFS Technology

M-Systems' products, which incorporate M-Systems' TrueFFS technology, use the following parameters to calculate worst case life expectancy:

- Media size (disk capacity)
- Flash disk endurance (write/erase cycles guaranteed by flash disk vendor)
- TrueFFS overhead factor: This factor represents the wear-leveling overhead 99.5% = 0.995
- Amount of data written per time unit (rate at which data is written)

The calculation is as follows:

$$LifeSpan(time) = \frac{MediaSize \times FlashDiskEndurance \times \mathrm{TrueFFSOverheadFactor}}{\frac{AmountOfDataWritten}{UnitOfTime}}$$

Examples

For the examples below, assume that:

- Media size (disk size) = 512MB (0.5GB)
- The application is writing data, filling the entire disk 240 times in a day (10 times every hour).
- Flash disk endurance:
  - "Erase before Write" = 1,000,000 write/erase cycles (today most flash disks vendors commit to between 300,000 to 2,000,000 write/erase cycles)
  - TrueFFS technology = 5,000,000 write erase cycles.

"Erase before Write" Algorithm

1. Worst Case Scenario
   The disk updates once per sector, and needs to erase an entire erasable block. This scenario is typical for updating the File Allocation Table (FAT):

   $$\text{Min Life Span} = \frac{512MB \times 1{,}000{,}000}{\frac{512MB * 240}{1day} \times 256} = \textbf{16 Days}$$

2. Best Case Scenario
   The disk updates once per full block (32 sectors), and needs to erase an entire erasable block. This scenario is typical of data acquisition systems that update large files.

   $$\text{Max Life Span} = \frac{512MB \times 1{,}000{,}000}{\frac{512MB * 240}{1day}} = \textbf{4,166 Days = 11 Years}$$

TrueFFS Technology

1. Best Case and Worst Case Scenario

   $$\text{Life Span} = \frac{512MB \times 5{,}000{,}000 \times 0.995}{\frac{512MB * 240}{1day}} = \textbf{20,729 Days = 56 Years}$$

CONCLUSION

The life expectancy of a flash disk depends in large part on the flash management tool used. The examples in this document demonstrate a significant improvement in flash disk endurance when TrueFFS technology is used. M-Systems' products have 5 to 1295 times better disk endurance than competing flash products.

M-Systems
Flash Disk Pioneers

HOW TO CONTACT US

USA
M-Systems Inc.
555 North Mathilda Avenue, Suite 220
Sunnyvale, CA 94085
Phone: +1-408-470-4440
Fax: +1-408-470-4470

China
M-Systems China Ltd.
Room 121-122
Bldg. 2, International Commerce & Exhibition Ctr.
Hong Hua Rd.
Futian Free Trade Zone
Shenzhen, China
Phone: +86-755-8348-5218
Fax: +86-755-8348-5418

Japan
M-Systems Japan Inc.
Asahi Seimei Gotanda Bldg., 3F
5-25-16 Higashi-Gotanda
Shinagawa-ku Tokyo, 141-0022
Phone: +81-3-5423-8101
Fax: +81-3-5423-8102

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

Taiwan
M-Systems Asia Ltd.
14 F, No. 6, Sec. 3
Minquan East Road
Taipei, Taiwan, 104
Tel: +886-2-2515-2522
Fax: +886-2-2515-2295

Internet
www.m-systems.com

General Information
info@m-systems.com

Sales and Technical Information
techsupport@m-systems.com