

From Wikipedia, the free encyclopedia

**Wear leveling** (also written **wear leveling**) is a technique<sup>[1]</sup> for prolonging the service life of some kinds of erasable computer storage media, such as flash memory.

The term has also been used by Western Digital to describe their hard disk preservation technique, but hard disks are not generally wear-levelled devices.

## Contents

- 1 Rationale
- 2 Techniques
- 3 References
- 4 External links

## Rationale

EEPROM and flash memory media have individually erasable segments, each of which can be put through a finite number of erase cycles before becoming unreliable. This can be anywhere between 10,000 and 1,000,000 cycles, for example, for NAND flash devices. Erasable optical media such as CD-RW and DVD-RW are rated at up to 1,000 cycles (100,000 cycles for DVD-RAM media).

Wear-leveling attempts to work around these limitations by arranging data so that erasures and re-writes are distributed evenly across the medium. In this way, no single sector prematurely fails due to a high concentration of write cycles.

Conventional file systems like FAT, ext2 and NTFS were originally designed for magnetic disks and as such rewrite many of their data structures (such as their directories) repeatedly in place. Some file systems aggravate the problem by tracking last-access times, which can lead to file metadata being constantly rewritten in-place.

## Techniques

There are several techniques for extending the life of the media:

- A checksum or error-correcting code can be kept for block or sector in order to detect errors or correct errors.
- A pool of reserve space can also be kept. When a block or sector does fail, future reads and writes to it can be redirected to a replacement in that pool.
- Blocks or sectors on the media can be tracked in a LRU (least recently used) queue of some sort. The data structures for the queue itself must either be stored off-device or in such a way that the space it uses is itself wear-levelled.

On flash memory devices, such as CompactFlash and Secure Digital cards, these techniques are implemented in hardware by a built-in microcontroller. On such devices, wear-leveling is transparent and most conventional file systems can be used as-is on them.

Wear-leveling can also be implemented in software by special-purpose file systems such as JFFS2 and YAFFS on flash media or UDF on optical media. All three are log-structured filesystems in that they treat their media as circular logs and write to them in sequential passes.

Some storage interfaces like DiskOnChip (which allows flash memory devices to emulate regular ATA disks) do not in themselves perform wear leveling.

## References

- ↑ U.S. Patent 6,850,443 (<http://patft.uspto.gov/netacgi/nph-Parser?patentnumber=6850443>) Wear leveling techniques for flash EEPROM systems.

## External links

- TrueFFS(R) Wear-Leveling Mechanism ([http://www.m-sys.com/NR/rdonlyres/FCC7D817-38A5-4D80-8471-67DA793EA255/0/TN\\_017\\_TrueFFS\\_Wear\\_Leveling\\_Mechanism.pdf](http://www.m-sys.com/NR/rdonlyres/FCC7D817-38A5-4D80-8471-67DA793EA255/0/TN_017_TrueFFS_Wear_Leveling_Mechanism.pdf)) – By Dmitry Shmidt (May 20, 2002) (M-Systems Technical Note TN-DOC-017 Rev. 1.1 91-SR-002-55-7L)
- Flash SSD wear-leveling and error-correction description by BitMicro. ([http://www.bitmicro.com/press\\_resources\\_flash\\_ssd.php](http://www.bitmicro.com/press_resources_flash_ssd.php))

Retrieved from "http://en.wikipedia.org/wiki/Wear\_leveling"

Categories: Articles with unsourced statements since July 2007 | All articles with unsourced statements | Computer memory | Non-volatile memory

- This page was last modified 16:43, 19 July 2007.
- All text is available under the terms of the GNU Free Documentation License. (See **Copyrights** for details.) Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a US-registered 501(c)(3) tax-deductible nonprofit charity.