swissbit®

Application Note

AN2110en

SSD Partitioning

[1] This work is licensed under the Creative Commons License "Attribution 4.0 International". To view a copy of this license, visit http://creativecommons.org/licenses/by/4.0/

swissbit®

Contents

1 Abstract

Partitioning can have a big influence on the speed and the lifetime of an SSD. Mainly with older industrial systems, whose operating systems still originate from the time before the wide spread of SSDs, this problem occurs frequently. The cause and remedies are discussed below.

2 Fundamentals

With hard disks, there is a fixed assignment between logical and physical addresses. If the host writes to the same logical block address (LBA) ten times, the same sector is addressed ten times under the same head in the same track, and its contents are replaced. With NAND flash, however, existing data cannot simply be overwritten because the area must be erased first. For technological reasons, erasing can only be done on entire flash blocks, which are typically several megabytes in size. Therefore, data that the host writes to the same logical block address is stored at a different position in the NAND flash each time. The memory position is thereby stored in a table also in the NAND, allowing quick access to the data during a read.

To manage the data in the NAND efficiently, manufacturers established a common size of 4 KiB. This is the smallest management unit used by today's file systems (so-called *cluster*).

3 Partitioning

It is important that the file system clusters are aligned with the management boundaries of the NAND flash. Otherwise, delays will occur, especially with random write or read accesses.

Figure 1 shows two examples of misaligned partitions. In each case, the partition is misaligned by one sector (512 bytes) against the administrative boundaries in flash. This often occurs when using an outdated partitioning tool designed only for hard disks. With hard disks, partitions started at the beginning of a track, and tracks had a maximum of 63 sectors. Therefore, the old partitioning programs historically set the start of a partition to sector number 63: the first track contained sectors 0 to 62; and the second track started with sector 63. This resulted in an offset of exactly one sector.

However, this offset should be avoided, if possible, since more 4 KiB units must be read during a read access than would be necessary with an aligned partition. In addition, wear-out increases because more data must also be written to the flash during a write access.

The partitioning tools of modern operating systems recognize SSDs and automatically align paritions correctly. These alignments occur not only at 4 KiB limits but mostly at 128 KiB or even 1 MiB to provide sufficient place for a possible boot loader. Only older systems, which are often found in the industrial environment, still use the old CHS (Cylinder/Head/Sector) scheme. In this case, a more modern tool should be used for the partitioning. A modern partitioning tool not only sets the start of the partitions to correct addresses, but it also automatically ensures the area for the files within the partition lies on such an address. With older tools, this often has to be done manually for FAT partitions by forcing reserved sectors between the boot sector and the first FAT table. However, changing the number of reserved sectors sometimes causes compatibility problems on systems running FAT16.

Such systems with outdated operating systems or old BIOS versions often allow a partition start only at a sector address divisible by

63. Then the resulting limitations in terms of wear-out and speed must be accepted. The use of single-level cell (SLC) flash usually offers a sufficient lifetime in such cases.

An exception are the Swissbit Longevity products, which do not support the modern 4 KiB management sizes. This affects the S-250, C-300L, C-350 and C-440 series. Here a partition should – if possible – be aligned to a 64 KiB boundary (e.g., start at sector 128, 256, or 384).

4 File system

Swissbit flash memories are basically compatible with any file system. However, the file system is of great importance regarding reliability in case of sudden power failures. While our products are protected against power failures – i.e., neither static data nor internal management data are damaged – file systems without a journal may experience major data loss. This primarily affects FAT file systems (FAT16 and FAT32), which are architecturally very sensitive to power failures when writing meta data. Therefore, wherever possible, only file systems with a journal, such as NTFS or ext4, should be used. These file systems keep a log before and after writing the meta data, which enables them to restore the consistency of the file system after a sudden power failure.

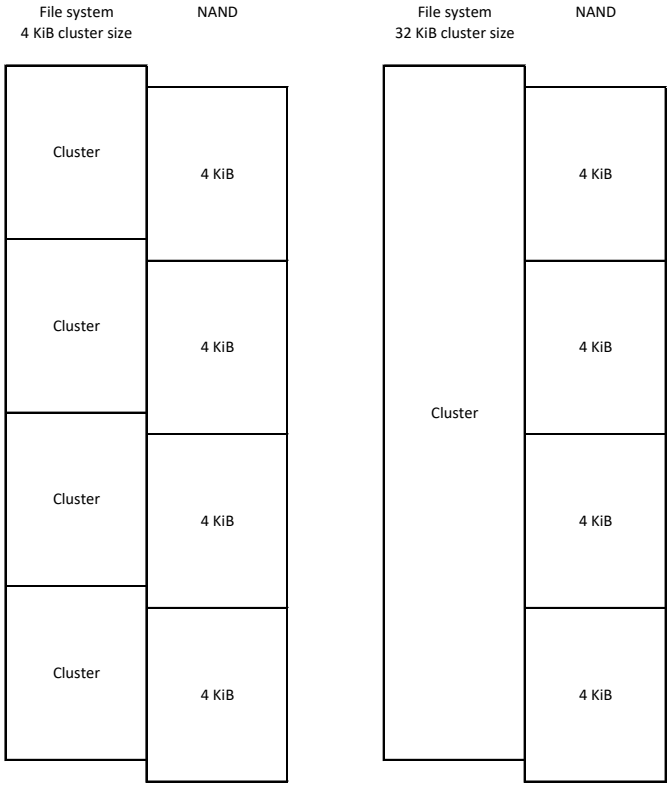

Figure 1: Misaligned partitions due to HDD partitioning program.
Left: 4 KiB cluster size (typical for NTFS and ext4)
Right: 16 KiB cluster size (32 GB FAT partition)

CONTACT US

Headquarters	**Swissbit AG** Industriestrasse 4 9552 Bronschhofen Switzerland	Tel. +41 71 913 03 03 sales@swissbit.com
Germany (Berlin)	**Swissbit Germany AG** Bitterfelder Strasse 22 12681 Berlin Germany	Tel. +49 30 936 954 0 sales@swissbit.com
Germany (Munich)	**Swissbit Germany AG** Leuchtenbergring 3 81677 Munich Germany	Tel. +49 30 936 954 400 sales@swissbit.com
North and South America	**Swissbit NA Inc.** 238 Littleton Road, Suite 202B Westford, MA 01886 USA	Tel. +1 978-490-3252 salesna@swissbit.com
Japan	**Swissbit Japan Co., Ltd.** CONCIERIA Tower West 2F 6-20-7 Nishishinjuku Shinjuku City, Tokyo 160-0023 Japan	Tel. +81 3 6258 0521 sales-japan@swissbit.com
Taiwan	**Swissbit Taiwan** 3F., No. 501, Sec.2, Tiding Blvd. Neihu District, Taipei City 114 Taiwan, R.O.C.	Tel. +886 912 059 197 salesasia@swissbit.com
China	**Swissbit China**	Tel. +886 958 922 333 salesasia@swissbit.com

Disclaimer:
The information in this document is subject to change without notice. Swissbit AG ("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.

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.