Challenges in Higher Education & Research, vol. 16
eds. T. Tashev, R. Deliyski, B. Lepadatescu, D. Hernandez-Dukova, Heron Press, Sofia, 2018

RESEARCH ON THE INFLUENCE OF THE USED OPERATING AND FILE SYSTEM ON THE TRANSFER RATE IN SSD DEVICES

Sergey Nedev[1], **Rumen Trifonov**[1], **Vladimir Kamenov**[2]

[1]*Technical University of Sofia, Faculty of Computer Systems and Technologies, 8 Kliment Ohridski, 1000 Sofia, Bulgaria*
[2]*Technical University of Sofia, Faculty of Mechanical Engineering, 8 Kliment Ohridski, 1000 Sofia, Bulgaria*

Abstract: In the presented article a study of SSD devices' transfer read/write rate dependence on the used file and operating system is made. Two SSDs are used to make the study more complete. The comparison was made under two operating systems - Windows 7 and Linux, using three file systems (four under Linux) and three software products for greater reliability, on the same computer. Also the characteristics of HDDs and SSDs, based upon different criteria, are compared. After the analysis of the results, conclusions and recommendations are made regarding the integration of SSD devices into computer configurations.

Keywords: hard drive, solid state drive, transfer rate, operating system, file system, NTFS, FAT 32, Ext4.

1. Introduction

SSD (solid state drive) is a static storage device that uses ICs. This technology primarily uses I/O interfaces compatible with traditional hard disks, allowing for easy HDD replacement with SSD. In addition, new interfaces are created to meet the requirements of SSD (SATA Express, for example). SSDs do not contain moving parts, which distinguishes them from traditional mechanical magnetic disks - hard disks (HDDs) or floppy disks containing rotating discs and moving heads performing I/O operations. Nowadays, most of the SSDs have a NAND flash memory (also used on USB flash drives) and are non-volatile. The SSD has two key components: controller and storage memory. The controller is a processor that performs instructions at a very basic level, making it extremely important for disk performance. The controller connects the NAND memory and the host computer. Some of the basic functions performed by the controller are: error correcting code (ECC); wear component management; marking bad sectors; caching of reading and writing and encryption and garbage collection.

NAND memories used to build SSDs can store one bit in each of their cells (SLC memory, two bits (MLC), or three bits (TLC). The latter provide the greatest density of the recording, respectively the larger volume of the device built through them. However, this is accompanied by a reduction in the number of memory overwriting cycles.

2. File System Used

SSD drives typically use the same file systems as hard drives. It is necessary for the file system to support the TRIM command, allowing the operating system to inform the static device which blocks of information are no longer considered usable and can be deleted. Some types of flash file systems (such as F2FS, JFFS2) help reduce the recording of information on the SSD, especially when only a very small amount of data is changed, for example, when updating the file system or processing metadata. The most common file systems that work well with the SSD are: Linux - file systems like ext4, Btrfs, XFS and JFS which includes TRIM technology. Since November 2013 ext4 can be recommended as a safe choice for a file system. F2FS is a state-of-the-art file system optimized for flash-based drives, and from a technical point of view it is a very good choice, but it is still in the experimental phase. In version 2.6.33 of the Linux kernel of February 24, 2010, support for the TRIM operation was introduced. Without support for TRIM, the SSD rate decreases over time. Mac OS X - versions after 10.6.8 (Snow Leopard) support TRIM, but only when used with an Apple SSD purchased device. There are technologies that enable TRIM activation in earlier versions of Mac OS X, although it is not certain whether TRIM works properly in versions prior to 10.6.8 [1]. Microsoft Windows versions prior to version 7 do not take special measures to manage SSDs. For all subsequent versions (7, 8, 8.1, 10), the NTFS standard file system supports TRIM. With these versions of the operating system, the TRIM commands are executed automatically when an SSD drive is detected and the operations are optimized. TRIM reduces the garbage collection function. Without TRIM support, the device will continue to record data during the garbage collection operation, which will further wear the disk. Windows disables defragmentation of solid state disks.

3. Characteristics of HDD and SSD

A comparison of the major features of hard disks and SSDs is shown in Table 1.

4. Experimental Layout

The tested SSDs are: Toshiba SSDQ300 (NAND 3 bit-per-cell, 2.5", SATA3 6 GB/s, 120 GB) and Transcend SSD370S (MLC NAND, 2.5", SATA3 6 GB/s, 128 GB). Research has been made on the same computer

Table 1. HDD and SSD comparison

Parameter	**SSD**	**HDD**
Start time	Almost instantly, there are no mechanical components that need to be started. It may take several milliseconds to exit automatic power saving mode	Rotating the discs takes a few seconds. Systems with many HDDs run them step-by-step to avoid high power consumption during start-up
Time for random access	Generally less than 0.1 ms, as data can be extracted directly from multiple locations on the flash memory. Access time is generally not a major performance limitation	Varies between 2.9 (high-performance server drives) and 2 milliseconds (laptop drives) due to head movement and waiting for data to fall under the read / write head
Delay of read time	Low because data can be read directly from anywhere. Faster start up occurs for programs that do a lot of searches on the hard drive	Much higher than the SSD. The rate is different for each search because the location of the data and the location of the read/write head is different
Transfer rate	SSD technology provides high read and write rates, but when accessing many individual and small information blocks, performance is reduced. The data transfer rate varies between 100 and 600 MB/s, depending on the disc. Up to several GB/s of data transfer rate for corporate products	When the head is positioned for read/write, the data transfer rate for a server disk is about 140 MB/s. In practice, this rate is much lower because of fragmentation and the search for data that is at different locations on the disk. The rate also depends on the rate of rotation of the discs (4,200 to 15,000RPM)
Reading	The read rate does not change depending on where the file is in the SSD. Unlike mechanical HDDs, SSD technology suffers from performance degradation, where NAND cells show a significant drop in performance and continue to degrade	When reading data is present on scattered locations on the disk, and when there is fragmentation, the response time increases considerably
Depending on the file system	There is little benefit in reading data sequentially, which makes fragmentation insignificant for SSD. Defragmentation would damage the SSD due to extra NAND subscriptions, which shortens the life of the SSD	Files of varying size over time gradually become fragmented if there are often write operations. In order to ensure good performance, it is necessary to periodically defragment the disk
Shock and vibration	There are no moving parts. Very resistant to shocks and vibrations	Heads that move freely over the disk are subject to vibration and can damage it
Acoustics	The SSD has no moving parts and is silent, except for electric noises that may arise from the electrical circuit	The HDDs have moving parts (head, actuator, motor) and give a distinctive sound of clapping and clicking. The sounds are different for different models. It is also considered that laptop HDD are much quieter
Temperature	Generally, the SSD does not need special cooling and can withstand higher temperatures	Temperatures above 35°C may reduce the life of the disc, as temperatures above 55°C may damage the disc. Here comes the need for properly designed cooling in the machine assembly
Influence of magnetic fields	The influence is small, but an electromagnetic pulse will damage any electrical system	In principle, the impact of magnetic fields or magnetic strikes can cause damage, even though disk drives are well protected with metal shields
Resilience and life	The SSD has no moving parts that can be mechanically damaged. Each block of information in the SSD can be erased and saved a number of times before it gets damaged. Controllers support this limit so that the units can last years in normal use. SSD-based on DRAM do not have this limitation, but if the controller fails, it can make the SSD unusable. Reliability varies considerably between different manufacturers. Many SSDs are critically damaged on power failure	Hard drives have moving parts and are susceptible to mechanical damage due to wear on the parts. The average life of a disc is 6 years. The likelihood of sudden and catastrophic data loss is lower for HDDs. Non-powered, hard drives hold the information longer than an SSD
Read/write performance	Cheaper SSDs generally have a write rate significantly lower than the read rate. Higher-quality SSDs have similar read and write rates	Hard disks require more time to search for read data than to write
Release of memory blocks	The writing capacity of the SSD is influenced by the free information blocks. Recorded data blocks that are not used cannot be processed by TRIM, however, even with TRIM, less free information blocks cause slower performance	Hard drives are not affected by free memory blocks and do not take advantage of the functionality of TRIM
Energy consumption	High-performance SSDs require 1/3 of the power of a hard drive (2-3watts)	2.5” HDD from 2 to 5W; 3.5” HDD - about 20 W (highest-performance, typically 7 W)

Lenovo G580 with processor Intel Core i5 - 3230M 2.6GHz (Turbo 3.2GHz); Memory: 6GB RAM DDR3, 1600MHz; OS: Ubuntu Linux 15.10, Windows 7 Ultimate SP1. Examined file systems are: Linux – FAT16; FAT32; NTFS; Ext4; Windows 7 – FAT16; FAT32; NTFS. The benchmark tools used are: Crystal Disk Mark 3.0.4 (portable, windows); Atto bench32 v.2.47 (windows); Gnome disks (Linux) [2] and [3].

Figure 1. Comparison of transfer rate for different file systems under Windows 7 (Toshiba SSD – top and Transcend SSD – below).

5. Study Results

The tests performed are summarized in Tables 2 and 3 as well as in Figures 1, 2 and 3. Average transfer read/write rates, under two operating systems (Linux and Windows 7), four file systems (Ext4, Fat16, FAT32 and NTFS) for two SSDs (Toshiba and Transcend) are shown.

Table 3. Transfer rates in Mb/s using different file systems under Windows 7

File System →	**NTFS**		**FAT16**		**FAT32**	
SSD Disc ↓	Read	Write	Read	Write	Read	Write
Transcend no cashing						
Overlapped I/O	406,4	157,5	396,1	155,9	416,3	157,6
I/O Comparison	368,1	156,8	355,1	156,7	357,0	157,3
Neither	389,2	154,1	394,0	156,0	396,0	158,0
Toshiba no cashing						
Overlapped I/O	430,2	88,6	418,8	91,0	427,3	86,5
I/O Comparison	361,3	88,4	358,7	89,2	359,6	86,6
Neither	395,9	92,9	400,5	88,6	400,4	90,0
Transcend with cashing						
Overlapped I/O	425,6	157,2	420,5	157,6	427,4	157,7
I/O Comparison	362,9	155,7	364,0	156,9	366,1	157,5
Neither	407,0	157,4	407,4	157,4	404,4	157,2
Toshiba with cashing						
Overlapped I/O	447,0	89,6	430,4	91,3	445,4	86,8
I/O Comparison	386,1	90,6	364,3	89,8	376,4	88,7
Neither	415,9	85,8	412,3	86,2	422,0	91,3

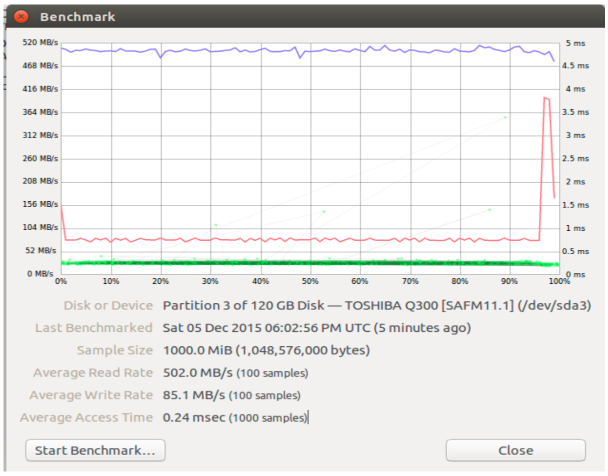

Figure 2. Comparison of transfer rate of Toshiba SSD under Linux Ext4 file system (average read rate - blue line, average write rate - red line, average access time - green points).

Table 2. Transfer rates in Mb/s using different file systems under Linux

SSD Disc	**NTFS**		**FAT16**		**FAT32**		**Ext4**	
	Read	Write	Read	Write	Read	Write	Read	Write
Transcend	503,4	70,2	508,2	124,5	517.8	98.7	516.5	135.8
Toshiba	499,9	110,1	496,0	165,9	496,4	173,8	502.0	85.1

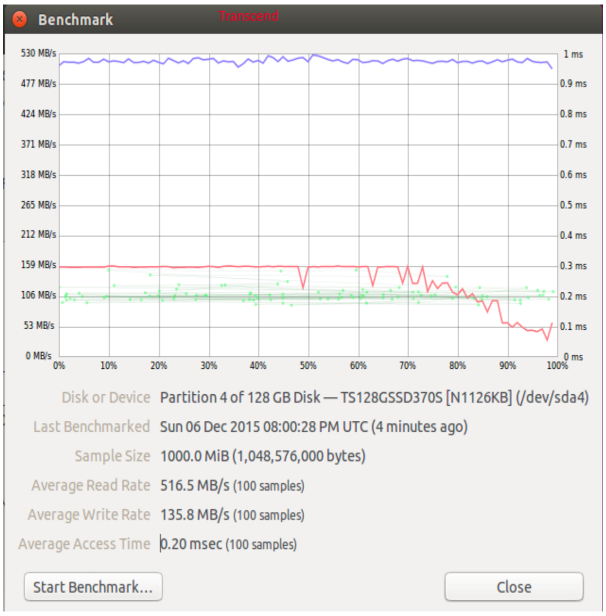

Figure 3. Comparison of transfer rate of Transcend SSD under Linux Ext4 file system (average read rate – blue line, average write rate – red line, average access time – green points).

6. Conclusions

When comparing the SSD manufacturers, under Windows 7 the Transcend drive is a little ahead of Toshiba drive, more pronounced at write rates. Under Linux, Toshiba once again leads in write rates.

When comparing the influence of the file system used - under Windows 7 (Figure 1) NTFS is ahead of other file systems both when writing and reading. Under Linux the Ext4 file system (Table 1) has a slight predominance over NTFS. Obviously, newer file systems show improved performance. The difference is not great, but at present 1-5 MB/s is a good data transfer rate, and this difference can be key when choosing a file system to use.

When comparing the influence of used operating system, it's obvious that under Linux SSDs are undoubtedly doing faster in reading than in Windows. When it comes to recording, however, it appears that there is a race between the two systems, but looking at the average values, it's clear that there is not so much difference in Windows, and all the studies have shown results in a narrower interval that is higher in values ??than under Linux. Of course, it should be borne in mind that different testing tools have been used, which may lead to discrepancies in the results.

In conclusion, the choose of a new SSD should be careful because the data that manufacturers publish for read and write times are the best they have achieved, but that does not mean that such a rate will always be reached in normal operation. A consideration of the type of intended operating and file systems the devices will be used on should be made. It's good to do a thorough analysis to avoid shortening the life of the selected SSD, as each operating system does not support SSDs without any additional settings or software (older OS). Compared to the HDD, SSD is a good investment, though with some worrying negative aspects, such as the life of the SSD compared to HDD.

References

[1] https://en.wikipedia.org/wiki/Solid-state_drive

[2] http://www.userbenchmark.com/ATTO-Disk-Benchmark/Rating/1285

[3] http://www.userbenchmark.com/Crystal-Dew-World-CrystalDiskMark/Rating/1284