Accepted: 8.09.2020. Original scientific paper

Reviewed: 12.11.2020. UDC: 004.3:004.71

Solid State Disk drive synthetic performances Analysis of 4th Gen. NVMe Protocol support

Analiza sintetičkih performansi Solid State Diska s podrškom za NVMe protokol

Aleksandar Skendžić

College of Information Technologies - VSiTe, Klaićeva 7, Zagreb, Croatia

e-mail: wireless82@gmail.com

Abstract: *This paper shows synthetic performance analysis of Solid State Disk drive that supports NVMe 4.0 protocol. Results are presented by using disk benchmarking tools Cristal Disk Benchmark and ATTO Disk tool on referent testing system. Also, synthetic tests were performed by measurement sequential read/write and random read/write performances with different queues depth and data block sizes of 4K, 32K, 256K and 8 MB. All results were compared with an older protocol standard NVMe 3.0 and also with SATA III standard.*

Keywords: *NVMe protocol, Solid State Drive (SSD), benchmarking*

Sažetak: *Ovaj rad prikazuje sintetičku analizu performansi Solid State Disk pogona koji podržava NVMe 4.0 protokol. Rezultati su predstavljeni upotrebom alata za usporedbu diska Cristal Disk Benchmark i ATTO Disk alata na referentnom sustavu za testiranje. Također, na testnom disku izvedeni su sintetički testovi mjerenjem sekvencijalnog čitanja / pisanja i slučajnog čitanja / pisanja s različitim dubinama redova i veličinama blokova podataka od 4K, 32K, 256K i 8 MB. Svi su rezultati uspoređeni sa starijim standardom protokola NVMe 3.0 kao i sa SATA III standardom.*

Ključne riječi: *NVMe protokol, Solid State Drive (SSD), benchmarking*

1. Introduction

Introducing Solid State Disk technology at consumer market in the 1976. and early 1990s was represented by the way of much faster data communication between central processing unit (CPU), Memory (RAM) and Disk drive (HDD) in regards to much more traditional electro-mechanical hard drives (Edwards, 2012). The main difference between these two technologies lays on completely different technology approach. While traditional hard drives technically are making primary of electric and mechanic components, SSD drives are based on flash memory chips. Second thing that is most important is communication bus which provides communication between CPU, memory controller (RAM) and I/O controller. Until SSD occurrence, the main data communication bus between CPU, memory and hard drive was parallel communication bus based at first on SCSI and PATA (IDE) standards from 1980s and after those much faster and newer serial SATA standards at the beginning of 2001s in its several versions: SATA I, SATA II and SATA III which came after and presents the newest and the latest version of standard (Guha et al., 2016). All mentioned communication bus standards above are operating at difference data speeds where it is worth that faster is better.

The newest communication standard is Non Volatile Memory Express (NVMe). NVMe is a communications standard developed specially for SSDs by a consortium of vendors including Intel, Samsung, Sandisk, Dell, and Seagate (Lacobi, 2019). The main benefits of new NVMe standard are hi data transfer speeds especially when SSD drive read/write performance is at stake. NVMe SSD flash array drives have low latency and high throughput which give high performance computing and real-time big data analysis (Jin, Ahn and Lee, 2018). Especially designed protocol was introduced by NVM Express Workgroup, which consists of more than 100 companies. One of the he first versions 1.1 was released at 2012.[1] There are several NVMe standard generations, so this paper will give a short review and test of the newest NVMe Gen 4. Technology and it will show performance measurement of tested SSD drive based on NVMe Gen 4. Protocol. The results will be comparable with the results of referential SSD drive based on NVMe Gen 3. Protocol. The testing platform was based on AMD Ryzen 9 3950X CPU with AMD X570[2] chipset. Results were obtained by Crystal Disk Mark testing software.

---

[1] Nvmexpress. About NVMe. https://nvmexpress.org/about/ (05.7.2020)
[2] AMD X570 chipset only fully supports NVMe Gen 4. Protocol. State at 5th May 2020.

2. About NVMe protocol

NVMe is high speed scalable storage protocol which is built on high speed PCIe lanes. It is designed for non-volatile memory media directly connected to Central Processing Unit (CPU) via PCIe interface. The main advantage of NVMe protocol is low latency data paths to the underlying media compared to legacy SAS or SATA protocols. Legacy IDE and SATA are based on Advanced Host Controller Interface (AHCI) technical standard. High level of comparable data between AHCI and NVMe are shown in Table 1.

NVMe architecture and performance ratings are based on multiple Input-Output (I/O) queues support (multiple deep queues). Also, NVMe supports up to 64K queues with each queue have 64K entries (65536 commands per queue). Creating maximum queues by NVMe host software is allowed by NVMe controller. Queues are designed such that I/O commands and responses to those commands operate on the same processor core and take advantage of the parallel processing capabilities of multi-core processors[3]. Legacy protocol such SAS and SATA can support single queues which support 256 commands (SAS) and 32 commands (SATA). Figure 1 shows connection of NVMe SSD drive which is directly connected to CPU via PCIe interface unlike SATA SSD drive which communicates with CPU over Host Bus Adapter (HBA).

Figure 1. *SSD via PCIe interface vs. I/O Controller / HBA connected SSDs*

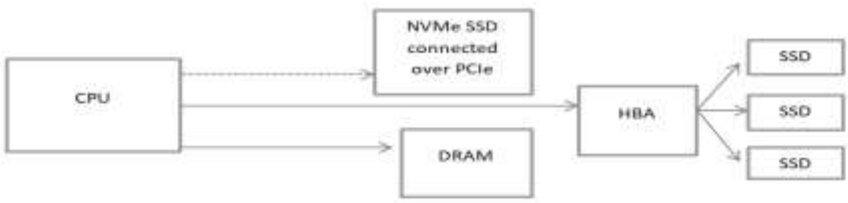

Source: (Gupta, 2020)

[3] NVMeOverview.https://www.nvmexpress.org/wpcontent/uploads/NVMe_Overview.pdf (05.7.2020)

Table 1. *High-level comparison of AHCI and NVMe*

Options	**AHCI**	**NVMe**
Maximum queue depth	One command queue; 32 commands per queue	65535 queues; 65536 commands per queue
Uncacheable register accesses (2000 cycles each)	Six per non-queued command; nine per queued command	Two per command
MSI-X and interrupt steering	A single interrupt; no steering	2048 MSI-X interrupts
Parallelism and multiple threads	Requires synchronization lock to issue a command	No locking
Efficiency for 4 KB commands	Command parameters require two serialized host DRAM fetches	Gets command parameters in one 64-byte fetch

Source: (Landsman, 2013.)

Table 1 summarizes high-level differences between the NVMe and AHCI logical device interfaces.

3. Synthetic vs. real tests

When we talk about benchmarking it is necessary to say that synthetic tests are different that real world benchmarks. Synthetic tests are presenting accurate comparison across different tests which are isolated but it may not reflect to user's experience. User experience comes with benchmarking complete system working together, how fast full system can complete task respectively. Hybrid benchmarking tests use combination of synthetic and real tests to get overall performance for completely test system.

SSD synthetic tests can measure read and write performance, both sequential and random tests, queues depth (threads) and compression. Data compression depends on SSD controller. The most controllers behave that they compress data before storing it on the SSD memory cells, so data that can be compressed can potentially be written and read quicker than uncompressible. Compression offers best case and worst case scenarios. Tests with compressed data offer worst case scenario, while tests with uncompressed data offer best case scenario.

SSD real world testing usually includes processing everyday tasks like software packages installing, file copying and archive extraction. Considering that SSD benchmarking is related on only one part of complete computer system, synthetic test is used.

4. Testing methodology, test configuration and results

Testing system was based on AMD CPU platform, 3rd generation of Ryzen 9 3900X which offers 12 CPU cores and 24 threads and which operates at base CPU clock of 3,8Ghz. Also, it has 64 MB of L3 Cache memory and supports PCIe 4.0 x 16[4]. The mainboard of the testing system was ASRock X570 Taichi. Chipset X570 is the only chipset that supports NVMe 4.0 at the moment of writing this text[5]. Tested SSD was GIGABYTE AORUS NVMe SSD of 2TB capacity which supports NVMe Gen 4. All synthetic tests where done by 64-bit version of **Crystal Disk Mark version 6.0.1** on Windows 10 64-bit platform and **ATTO Disk Benchmark version 4.01** which measured transfer data rates of specific lengths, for test usually 0.5 to 8192 kilobytes.

Testing results are based on several tests as follows:

- Sequential Read/Write with multi Queues and Threads (Seq Q32T1);
- Random 4KiB Read/Write with multi Queues and Threads (4K Q32T1 and 4K Q8T8);
- Sequential (Block Size=1MiB) Read/Write with single Thread (Seq);
- Random 4KiB Read Write with single Queue and Thread (4K Q1T1).

Queue Depth and number of Threads that are using are relating on multiple read and write processing at same time, which means that one input/output operation does not need to complete before the next one can start. SSD controller transfers data with multiple memory cells at same time and also higher queue depths will perform better performance with multitasking and multithreaded batch processing.

Crystal Disk Mark for test uses random data which is uncompressible. Uncompressible data presents a worst case scenario for SSD drive, such for storing images files or other data types that cannot be easily compressed (JPG files for example)[6]. ATTO disk benchmark for test is using different transfer sizes and at queue depth of 4 by default. This fits four threads simultaneously accessing the drive. While Crystal Disk Mark operates with uncompressible

[4] AMD Ryzen 9 3900X. http://www.amd.com/en/products/cpu/amd-ryzen-9-3900x/ (05.7.2020)

[5] NVMe 4.0 is fully supported only on AMD x570 chipset at the moment. The newest AMD B550 chipset supports NVMe 4.0 only on CPU graphics and storage. NVMe 4.0 is not backward compatible with older mainboard chipsets such AMD x470.

[6] Guide to SSD reviews and real World vs. synthetic tests. https://club.myce.com/t/a-guide-to-ssd-reviews-and-real-world-vs-synthetic-tests/398979 (07.7.2020)

data, ATTO disk benchmark operates with uncompressed data with the best case scenario and with the data being highly compressible. ATTO test was run with the default of 0.5KB and 64MB transfer sizes and 256MB of total length.

Generic performance of NVMe Gen 4.0 (PCIe 4.0) is offering data transfer rates of 5 GB/s or above while PCIe Gen 3.0 x4 is offering lower speeds such 4 GB/s. For example, the fastest SATA III standard is offering 550/520MB/s sequential read and sequential write speed rates respectively.

Figure 2. *Crystal Disk Mark - Test results of tested system based on NVMe Gen 4.0 protocol*

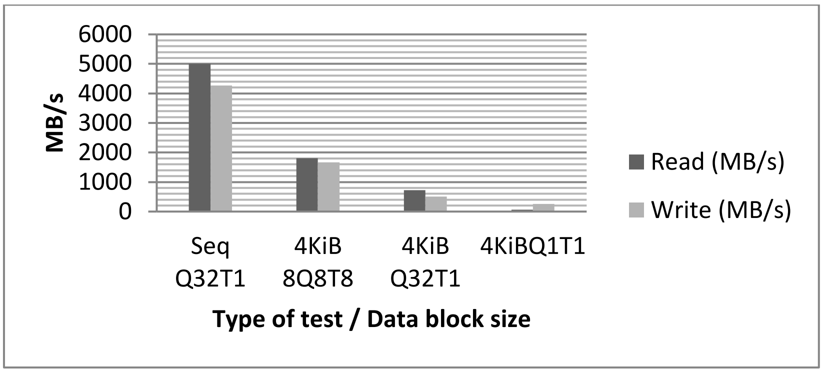

Source: author

As shown in Figure 2 sequential read/write performances with 32 queues depth and used one CPU thread amounts **5000.2 MB/s** of read and **4264 MB/s** of write. These synthetic results were expected, especially Read performance according to NVMe 4.0 specifications. Comparable synthetic test results of NVMe 3.0 are **3490 MB/s**. of read and **3466 MB/s** of write. These synthetic tests presents **69.79 %** better synthetic performance of NVMe Gen 4. in read and nearly **81.28 %** in write than NVMe Gen 3. Note that random 4KQ1T1 tests are the slowest and that refers to the device effectiveness, how quickly device can retrieving small parts of data from random locations respectively) (Arunović, 2020). That is very demanding for SSD controller because of the need to prepare (read), delete and write at cells many times in small blocks.

Figure 3. *Crystal Disk Mark - Test results of tested system based on NVMe Gen 3.0 protocol*

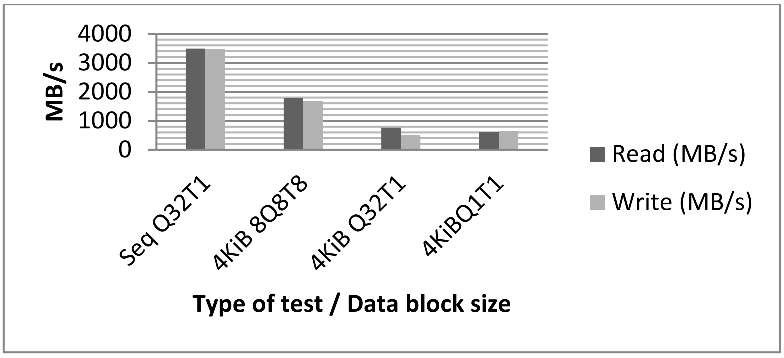

Source: author

While NVMe Gen 4.0 is offering sequential transfer rates reaching 5,000 MB/s or more, PCIe gen 3.0 x4, offers theoretical limit of 4 GB/s when a x4 interface is used. As shown in Figure 3. synthetic Read/Write tests showed that these results are much closer to **3500 MB/s**.

Figure 4. *ATTO Disk Benchmark - Test results of tested system based on NVMe Gen 4.0 protocol*

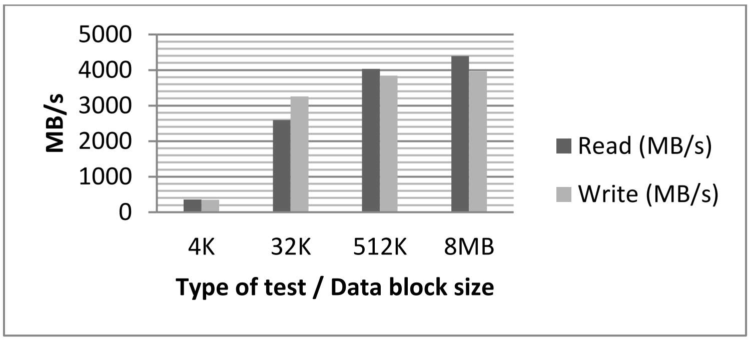

Source: author

During the test, tested system SSD showed some slower performance using the smallest data block of 4K where measured data were **361.07 MB/s** (Read) and **348.28 MB/s** (Write). When block was increased to the size of 32K, Read/Write performances were much higher such measured: **2585 MB/s** (Read) and **3256 MB/s**. For the test where block size were 512K and 8MB performance were also higher. For showing significant SSD performance improvement of NVMe 4.0, also some SATA III SSD test results were compared.

For the referent SATA III drive Samsung SSD 860 2TB drive is used. In synthetic tests in Crystal Disk Mark of tested Samsung reached 561MB/s read and 530 MB/s write in the sequential write test (Q32T1). Random 4K (4KQ1T1) test was 42.09 MB/s read and 108.4 MB/s write. Also note that 4K tests are much slower that sequential (32K) tests.

Figure 5. *The results of synthetic test of the SATA III Samsung 860 2TB SSD drive*

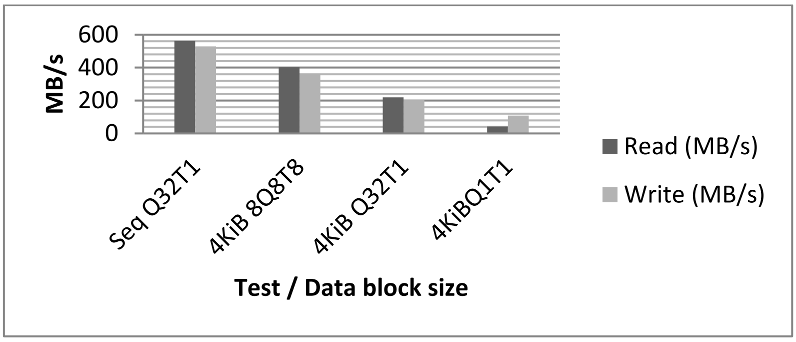

Source: author

All tests were performed such sequential read/write using 32K queues and random read/write using 4K data block size. Maximum results were 561.8 MB/s in sequential read and 530.1 MB/S in write. Random 4K read was 42.9 MB/s and write was 108.4 MB/s. Sequential read/write results are expected and comparable to SSD specifications and maximum SATA III speed of the manufacturer.

Figure 6. *Comparison of Data transfer speed SATA III vs. NVMe Protocol*

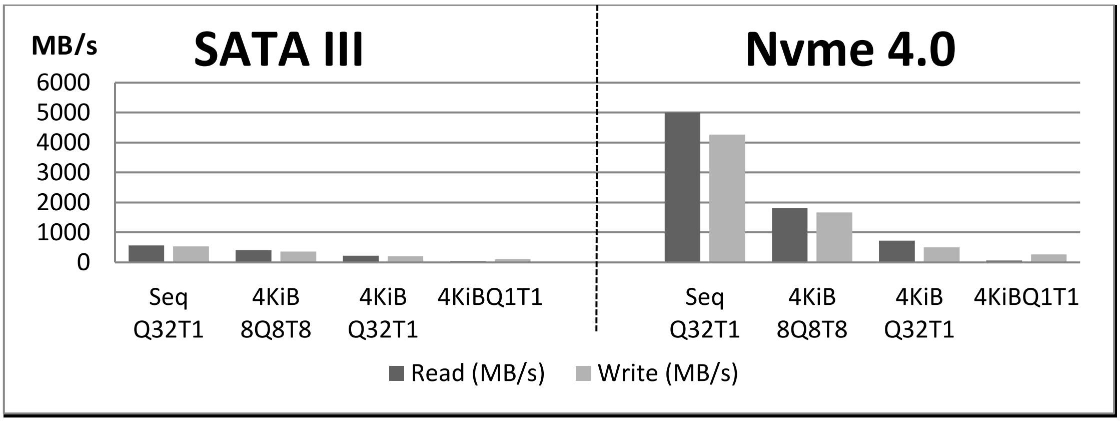

Source: author

As shown in Figure 6. NVMe 4.0 has significant improvement in sequential Read which is **9 times** faster than SATA III, 5000.2 MB/s vs 568.8 MB/s. On the other side random 4K tests which operate with small parts of data from random locations showed lower difference results but still significant. Read results of 4K tests showed slightly **over 50 %** of improvement (42.9 MB/s vs. 62.5 MB/s) in read and **more than 100 %** improvement in write (108.4 MB/s vs.

262.8 MB/s). Nearly 20% of overall disk access on a PC will consist of random reads[7]. Note significant difference in sequential Read/Write tests Q32T1.

5. Conclusion

These synthetic tests on tested platform and SSD which were performed by two popular Disk utility programs such Crystal Disk Benchmark and ATTO Disk benchmarking tool were showed relevant performance improvements of NVMe 4.0 regards to NVMe 3.0 protocol. Tests were performed as Sequential Read/Write with multi queues and threads and random block sizes (4K) with single queues and multi threads. Also, with ATTO disk tool test were performed with data block size of 512K and even 8MB. Performed synthetic test results showed that NVMe 4.0 has 69.79 % better synthetic performance in read and nearly 81.28 % in write than NVMe Gen 3.

These synthetic test confirmed significant improvement with 4th Gen of NVMe protocol especially when compared with an old SATA III. It is worth to know that SSD drives have lower performance when working with large amounts of small size files and when working with compressed data (Skendžić, Kovačić and Tijan, 2016). Also, the performance starts to degrade when SSD reaches around half of its capacity. For future research/tests authors recommend to prepare and run hybrid benchmarking tests which use combination of synthetic and real tests to get overall performance for completely test system.

References

1. AMD Ryzen 9 3900X. http://www.amd.com/en/products/cpu/amd-ryzen-9-3900x/ (05.7.2020)
2. Arunović, D. (2020). AMD Ryzen recenzije Ryzen 3100 i 3300x. https://www.bug.hr/recenzije/amd-ryzen-3-3100-i-3300x-brzi-i-pristupacni-15039 (07.7.2020)
3. Edwards, B. (2012). Evolution of the Solid State Drive. https://www.pcworld.com/article/246617/evolution-of-the-solid-state-drive.html (02.7.2020)
4. Guide to SSD reviews and real World vs synthetic tests. https://club.myce.com/t/a-guide-to-ssd-reviews-and-real-world-vs-synthetic-tests/398979 (07.7.2020)

[7] Userbenchmark (2020). https://www.userbenchmark.com/Faq/What-is-4K-random-read-speed/28 (07.7.2020)

5. Gupta, R. (2020). NVMe important data driven businesses. https://blog.westerndigital.com/nvme-important-data-driven-businesses/ (09.7.2020)
6. Y. T. Jin, S. Ahn and S. Lee, (2018). Performance Analysis of NVMe SSD-Based All-flash Array Systems. *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Belfast, 2018, 12-21.
7. Lacobi, L. J. (2019). Everything you need to know about NVMe. https://www.pcworld.com/article/2899351/everything-you-need-to-know-about-nvme.html (02.7.2020)
8. Landsman, D. (2013). AHCI and NVMe as Interfaces for SATA Express Devices – Overview. SATA - International organization.
9. NVMeOverview.https://www.nvmexpress.org/wpcontent/uploads/NVMe_Overview.pdf (05.7.2020)
10. Nvmexpress. About NVMe. https://nvmexpress.org/about/ (05.7.2020)
11. S. Guha; Wen Wang, S.; Ibraheem, M. Balakrishnan and J. Szefer, (2016). Design and implementation of open-source SATA III core for Stratix V FPGAs. *International Conference on Field-Programmable Technology (FPT)*, Xi'an, 2016, 237-240.
12. Skendžić, A.; Kovačić, B.; Tijan, E. (2016). Effectiveness analysis of using Solid State Disk technology. 39th International Convention – MIPRO**,** 1587-1591.
13. Userbenchmark (2020). https://www.userbenchmark.com/Faq/What-is-4K-random-read-speed/28 (07.7.2020)