ISSN(Online): 2320-9801
ISSN (Print): 2320-9798

International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Website: www.ijircce.com

Vol. 5, Issue 3, March 2017

Performance Evaluation of SATA and NVMe Based SSDs

Gitesh Hulke, Prof. Varshapriya J. N.

M. Tech. (NIMS) Student, Department of CE & IT, VJTI, Mumbai, India

Assistant Professor, Department of CE & IT, VJTI, Mumbai, India

ABSTRACT: With the increase in importance and dependence on information, the field of Information Storage has witnessed tremendous growth. Not very long ago, Hard Disk Drives (HDD) were used as the main storage devices. These HDDs are a combination of electrical and mechanical parts and usually rely on SATA/AHCI for storage and retrieval of data. Nowadays, however, there has been a gradual shift towards a different storage type called the Solid State Drives (SSD). The SSDs are built around the flash storage technology and do not have any moving parts like HDDs did. They can thus be governed by the various different interfaces and protocols which also include SATA, which was primarily developed for spinning disk HDDs. This paper discusses various protocols and interfaces used in SSDs and explains why NVMe might be the future of SSDs.

KEYWORDS: Solid State Drives, SATA, NVMe, PCI-Express

I. Introduction

Sincethe turn of this century, processors speeds have continued to increase dramatically while read and write times for mechanical hard disk drives (HDDs) have not. Today's CPUs can process data much faster than HDD storage can supply it, resulting in latency.

Solid State Drives (SSDs) bridge this gap to a large extent. SSD devices make use of silicon-based memory chips as the storage media for the writing and reading of persistent data. Unlike a spindled hard disk drive (HDD), an SSD does not contain any mechanical parts. These SSDs are available in a wide range of form factors depending on the requirements of the systems in which they are placed. In order to support compatibility with the older systems, some SSDs are designed with the standard SATA or SCSI interfaces. Although, these interfaces were primarily designed for HDDs, they can still be used for SSDs. The performance of such SSDs differs from the ones based on interfaces and protocols which are built around the flash memory based storages.

II. Related Work

High speed interconnect like PCIe enable fast communication between devices. PCIe enables devices to be directly connected to the host motherboard thereby minimizing latency. SSDs today are being designed to utilize these PCIe capabilities. In order to do so, a different set of interface and protocols need to be employed. NVM Express (NVMe) or Non-Volatile Memory Host Controller Interface Specification (NVMHCI) is a logical device interface specification for accessing non-volatile storage media attached via a PCI Express (PCIe) bus [1].

Another widely used standard for communication between storage devices and the host is SATA. Most consumer grade storage devices today make use of it and technologies and protocols associated with them. SATA is well suited for hard disk storages and in order to harness the speeds that the solid state disks are capable of providing, a different standard like the one mentioned above (NVMe) are implemented.

ISSN(Online): 2320-9801
ISSN (Print): 2320-9798

International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Website: www.ijircce.com

Vol. 5, Issue 3, March 2017

III. Solid State Drive

A solid-state disk or solid-state drive (SSD) is a type of non-volatile storage device which is used for storing persistent data on solid-state flash memory. As opposed to a traditional HDD which consists of a spinning disk with a read/write head on a mechanical arm, known as an actuator. The SSD, has an array of semiconductor memory organized as a disk drive, using integrated circuits rather than magnetic or optical storage media. There are two types of flash memory architecture:

A. NAND Flash

A NAND based flash consists of transistors which are connected in series. NAND Flash offers a maximum of million erase cycles per erase block.

B. NOR Flash

A NOR based flash consists of transistors which are connected in parallel and it offers a 10,000 to 100,000 erase cycles per erase block.

Due to better endurance as compare to NOR Flash, the NAND flash is more predominantly used. NAND devices offer up to 10 times the life span of NOR devices [2]. There are four types of NAND flash, differing in number of Program/Erase cycles per lifetime, and defined by their construction:

1) Single Level Cell (SLC)

   The charge stored in SLC may be interpreted either as 0 or 1. This results in a higher per unit storage cost since less data stored per cell. It is therefore more expensive and generally the fastest. It also has higher endurance and longer life.

2) Multi-Level Cell (MLC)

   The charge stored in MLC may be interpreted as a variety of values, 0 to 3, i.e. 4 possible states (2 bits). It has a shorter lifetimes which is usually 10x less than SLC. The advantage of this memory is that the cost is two to four times less than that of the SLC. The MLC is not as fast as the SLC and typically uses some form of error correction code per block.

IV. SATA Based SSDs

Solid-state drives have traditionally been designed to use the SATA interface to connect storage to servers, using host bus adapters and other components.

A SATA (Serial - Advanced Technology Attachment) is a serial version of the IDE/ATA (Integrated Device Electronics) specification. SATA is a disk-interface technology that was developed by a group of the industry's leading vendors. Serial ATA is a high-speed serial link replacement for the parallel ATA attachment of mass storage devices. The serial link employed is a high-speed differential layer that utilizes Gigabit technology and 8b/10b encoding [3].

It provides point-to-point connectivity for short distance of up to one meter and enables data transfer at a speed of 150 MB/s. Enhancements to the SATA have increased the data transfer speed up to 600 MB/s. [4]

A SATA bus directly connects each storage device to the host through a dedicated link, making use of low-voltage differential signaling. SATA devices are hot-pluggable, which means that they can be connected or removed while the host is up and running. A SATA port permits single-device connectivity.

SATA based SSDs are widely used because they are almost always backward compatible and low cost. The underlying command set is similar to the one employed for hard disk drives.

V. NVMe Based SSDs

Modern server-based flash storage involve SSDs that are designed for installation in the PCIe (Peripheral Component Interconnect Express) slots. These PCIe enabled SSDs are directly attached to the motherboard and

ISSN(Online): 2320-9801
ISSN (Print): 2320-9798

International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Website: www.ijircce.com

Vol. 5, Issue 3, March 2017

communicate with the server using a dedicated point to point connection. This eliminates the resource contention and reduces latency.

PCI Express is a high performance, general purpose I/O interconnect defined for a wide variety of future computing and communication platforms. Key PCI attributes, such as its usage model, load-store architecture, and software interfaces, are maintained, whereas its parallel bus implementation is 5 replaced by a highly scalable, fully serial interface. PCI Express takes advantage of recent advances in point-to-point interconnects, Switch-based technology, and packetized protocol to deliver new levels of performance and features [5].

PCIe SSD come is different form factors and all of them require some protocol in order to communicate with the host. SSD vendors are developing PCIe devices around the emerging Nonvolatile Memory Express (NVMe) protocol, a set of specifications that are designed to operate at the host controller level. NVM Express specifications are developed by an industry consortium, NVM Express, Inc. NVMe specifications aim to provide a more efficient I/O stack that takes full advantage of semiconductor flash storage technology and increases the throughput of PCIe devices.

NVM Express (NVMe) is a register level interface that allows host software to communicate with a non-volatile memory subsystem. This interface is optimized for Enterprise and Client solid state drives, typically attached to the PCI Express interface [1].

Reference [1] shows, the following key attributes of the interface:

- Does not require un-cacheable / MMIO register reads in the command submission or completion path.
- A maximum of one MMIO register write is necessary in the command submission path.
- Support for up to 65,535 I/O queues, with each I/O queue supporting up to 64K outstanding commands.
- Priority associated with each I/O queue with well-defined arbitration mechanism.
- All information to complete a 4KB read request is included in the 64B command itself, ensuring efficient small I/O operation.
- Efficient and streamlined command set.
- Support for MSI/MSI-X and interrupt aggregation.
- Support for multiple namespaces.
- Efficient support for I/O virtualization architectures like SR-IOV.
- Robust error reporting and management capabilities.
- Support for multi-path I/O and namespace sharing.

Note: The terms NVMe based SSD and PCIe based SSD are interchangeably used and refer to the same type of device.

VI. COMPARISON BETWEEN SATA-BASED SSD AND NVME-BASED SSDS

Since SATA based SSDs are compatible with even the older motherboards, installing a SATA based SSD does not require additional work. The PCIe based SSD on the other hand requires support from the BIOS. Although some old motherboards may not support PCIe SSD and may require a BIOS update, but, almost all modern motherboards natively support PCIe based SSDs.

SATA and PCIe based SSD defer in the underlying protocols they use. In order to compare the two types of SSDs, the same workstation setup was used as the host machine. For benchmarking the SSDs, CrystalDiskMark was used. CrystalDiskMark is a disk benchmark utility that measures performance for sequential and random reads/writes of various sizes for any storage device. It can measure sequential reads/writes speed, measure random 512 KB, 4 KB and 4 KB (Queue Depth = 32) reads/writes speed. The SATA based SSD used in the experiment was Samsung 850 Pro with 1TB of storage space, while the PCIe based SSD was Samsung 950 Pro with 1.2TB of storage. The test was carried with a test file workload of 1000MB i.e. 1GB approximately with 5 tests and their average was calculated and taken into consideration. The benchmarking tool test the SSD for sequential read/write as well as random read/write. 4K QD32 refers to random read/write with 32 as the depth of I/O queue of the SSD.

ISSN(Online): 2320-9801
ISSN (Print): 2320-9798

International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Website: www.ijircce.com

Vol. 5, Issue 3, March 2017

Table 1. SATA & NVMe SSD Performance comparison

Characteristic	**Solid State Disk**	
	Samsung 850 PRO	Samsung 950 PRO
Memory Size	1TB	1.2 TB
Interface	SATA 6Gb/s	PCIe 3.0 NVMe
Sequential Read Speed	533 MB/s	1610 MB/s
Sequential Write Speed	486 MB/s	1307 MB/s
512K Random Read Speed	403 MB/s	1101 MB/s
512K Random Write Speed	464MB/s	1296 MB/s
4K Random Read Speed	37 MB/s	35 MB/s
4K Random Write Speed	122 MB/s	214 MB/s
4K QD32 Random Read Speed	372 MB/s	711 MB/s
4K QD32 Random Write Speed	320 MB/s	638 MB/s

Table I. shows the performance comparison of SATA based SSD with the NVMe based SSD. It is evident from the results that the NVMe based SSD fared better than SATA based SSD in every attribute ranging from sequential read to 4K Queue Depth 32 writes.The sequential read speed of NVMe based SSD was three times more than that of the SATA based SSD. In most cases, the NVMe based SSD fared at least twice as fast as it SATA counterpart. This is because NVMe takes full advantage of the semiconductor technology by modifying the I/O stack to be more suitable for devices with flash storage.

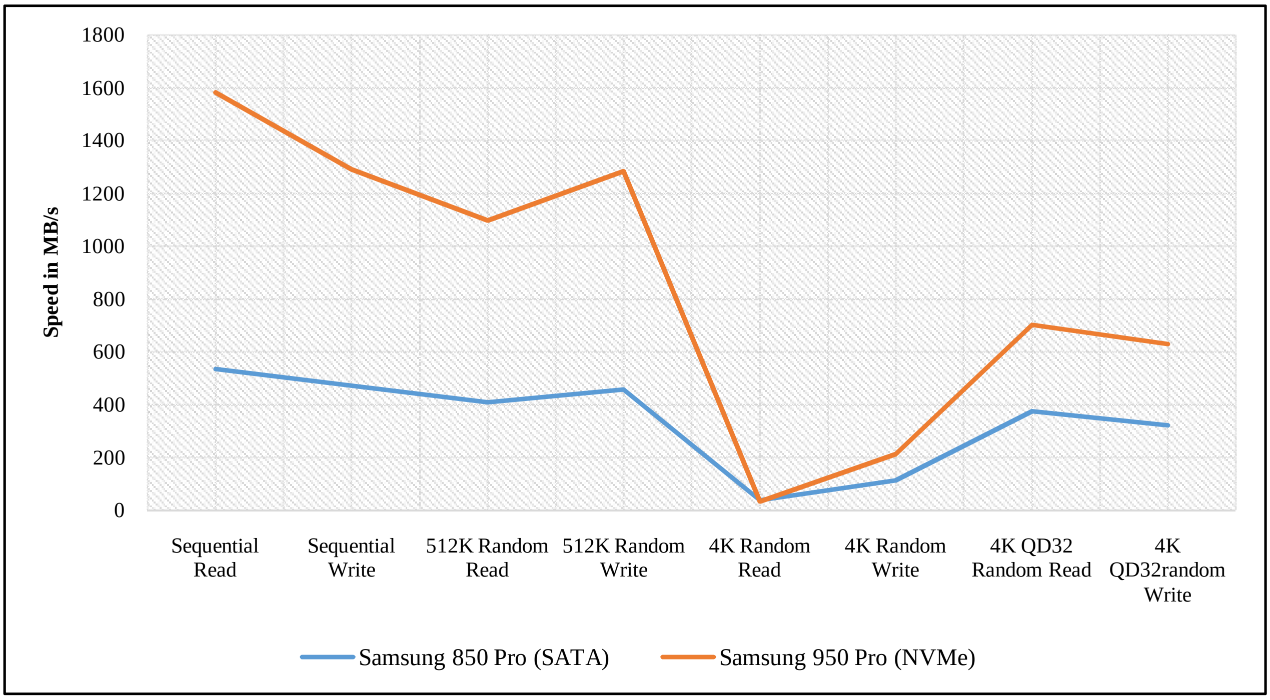

Figure 1. SATA based SSD Vs NVMe based SSD

ISSN(Online): 2320-9801
ISSN (Print): 2320-9798

International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Website: www.ijircce.com

Vol. 5, Issue 3, March 2017

Figure 1 depicts a graph which compares the Samsung 850 Pro SATA SSD to the NVMe based Samsung 950 Pro. Apart from 4K random reads, NVMe based SSD fares better than the SATA based one.

The NVMe based SSDs can have multiple I/O queues. Theoretically, queues as many as 65,535 are possible, with each I/O queue supporting up to 64K outstanding commands. This improves throughput which greatly impacts the performance. The SATA, on the other hand, acts as a bottleneck. Although the SSD is capable of higher read/write speeds, it is limited by the SATAs' design, which does not harness its full potential.Another aspect of the comparison is that of the cost. The NVMe based SSDs cost more than the SATA based ones. In some cases, the cost of an NVMe based SSD is twice or thrice of that of SATA based SSDs.

VII. Conclusion and Future Work

The SATA based SSDs although an upgrade from the HDDs are no match for the NVMe based ones. The SATA based SSDs are at least five times the speed of the traditional HDD. As compared to the HDDs, SSDs are more durable and produce less noise. The power consumption of an SSD is very less as compared to the HDD, making it suitable for devices that operate on batteries. SSDs also produce less heat and are shock resistant.The NVMe-based SSD can be up to four times faster than a SATA based SSD and therefore almost ten times faster than the traditional HDD. The result is increased performance. The higher performance of NVMe based SSD makes it particularly suitable for buffering and caching applications, with content delivery high on the list of suitable applications. NVMe based SSD is also a suitable technology for loading databases to significantly increase performance

Taking into consideration the factor of cost associated with each of the above mentioned storage devices, the utility of each device varies. In scenarios where speed is not as big a factor as cost, the use of HDD will suffice. In scenarios where the user need to load or process large files quite frequently, the SATA based SSD can prove to be very useful. While the NVMe based SSD is more suitable for cloud and enterprise environments where cost of procurement does not matter as much as the performance. It can be used in cases that require high responsiveness under heavy workload.

References

1. NVM Express Inc. (Nov. 3, 2014), NVM-Express Specification Revision 1.2, [Online]. Available: http://www.nvmexpress.org/wp-content/uploads/NVM_Express_1_2_1_Gold_20160603.pdf. [Accessed: Jan. 23, 2017].
2. Arie Tal (Oct. 2002), Two Flash Technologies Compared: NOR vs NAND, [Online]. Available: https://focus.ti.com/pdfs/omap/diskonchipvsnor.pdf. [Accessed: Jan 23, 2017]
3. Serial ATA International Organization (June. 2, 2009Serial ATA Revision 3.0, [Online]. Available: http://www.lttconn.com/res/lttconn/pdres/201005/20100521170123066.pdf. [Accessed: Jan. 20, 2017]
4. G. Somasundaram, Alok Shrivastava, EMC Education Services,"Information Storage and Management: Storing, Managing, and Protecting Digital Information", 2nd ed., Wiley Publishing Inc., 2009, pp. 102-103.
5. PCI-SIG. (Nov. 10, 2010), PCI-Express Specification, [Online]. Availablehttps://http://composter.com.ua/documents/PCI_Express_Base_Specification_Revision_3.0.pdf. [Accessed: Jan. 20, 2017].
6. NVM Express Inc. (Dec. 28, 2014), NVM-Express Explained Whitepaper, [Online]. Availablehttp://www.nvmexpress.org/wp-content/uploads/2013/04/NVM_whitepaper.pdf. [Accessed: Jan. 23, 2017].
7. Y. Luo, Y. Cai, S. Ghose, J. Choi and O. Mutlu, "WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management," in Mass Storage Systems and Technologies (MSST), 2015.