Solid State Drive Architecture

A comparison and evaluation of data storage mediums

Tyler Thierolf
Justin Uriarte

Outline

- Introduction
  - Storage Device as Limiting Factor
- Terminology
  - Internals
- Interface
- Architecture
- Advantages
- Disadvantages
- Conclusions

Introduction

- To gain a full understanding of solid state drives (SSDs), one must be knowledgeable of the current generation of platter based drives
- SSDs mirror the functionality of the existing standard of hard disk drives (HDDs)

Storage Devices as Limiting Factor

- CPU Performance has improved approximately 60% every year
- I/O Sub-System Performance has improved less than 10% each year due to limitations imposed by the mechanical components
- Amdahl’s Law dictates that the disk I/O bottleneck diminishes the improvements introduced by increased RAM and CPU speed to the overall system performance
  – “Only as strong as your weakest link”

HDD vs. SSD

- Same form factors
  - 1.8’’, 2.5’’, and 3.5’’
- Interface
  - SATA, PCIe, SCSI
  - Computer doesn’t necessarily know which is connected

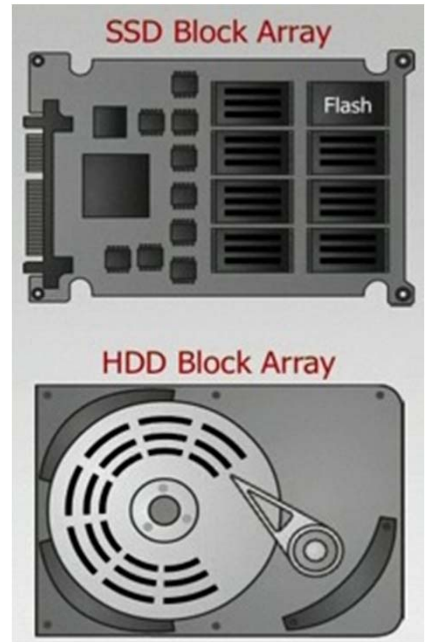

HDD

- Rotating Magnetic Media
  - On a disk called a platter
  - Platter rotates several hundred times a second
    - 5400, 7200, 10000, 15000 rpm
  - Platter contains magnetic domains on which data is written

HDD

- Always contain more than one platter
- There is a head for each platter, but the heads cannot move independently
- Cylinder – All tracks accessible without moving the head assembly
- Head – The device that writes data to and reads the surface of one side of a platter

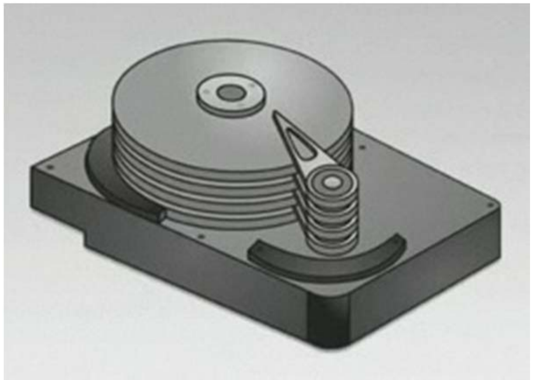

HDD Data Structure

- Sector – A slice of the platter that contains the minimum addressable read/write portion
  - Typically 512 bytes

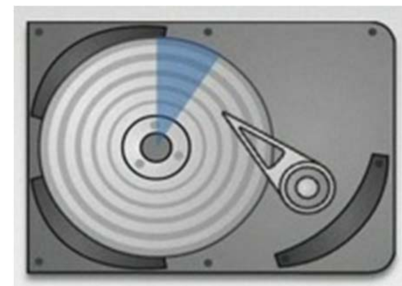

- Track – Thin concentric circular strips that contain portions of multiple sectors
  - Head can access all data on a track without moving

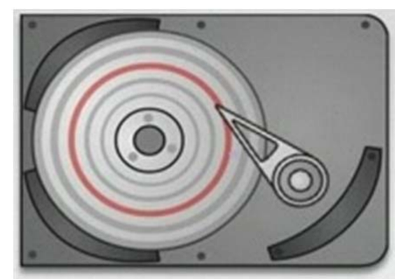

HDD Data Structure

- Blocks – The intersection of a track and a sector
  - Minimum addressable size in the HDD
  - Addresses are specified by providing the cylinder, head, and sector number

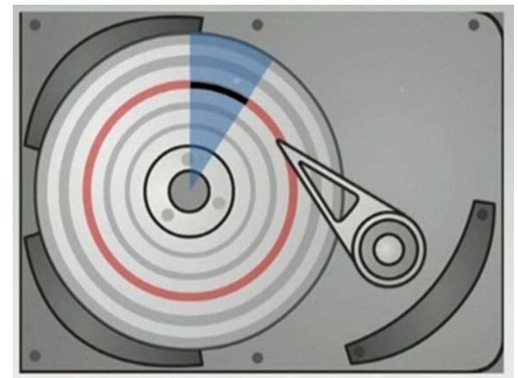

Logical Block Addressing (LBA)

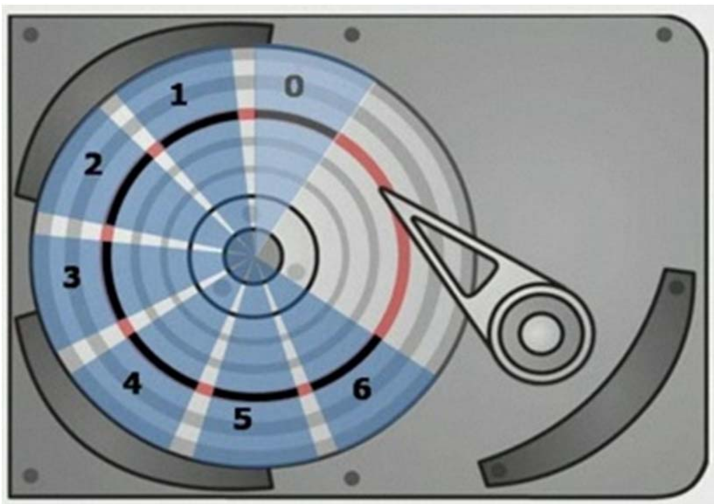

- A way of addressing blocks that simply numbers them linearly rather than providing a cylinder, head, and sector number
- This scheme is generally replacing the legacy block addressing scheme although both are supported on current SSDs and HDDs

HDD vs. SSD

- Sectors of concentric circles in HDDs represent the LBAs in the system
- In an SSD the LBAs are actually inside of the flash media
  - They represent individual addresses

Interface

- Serial ATA (SATA)
  - Serial: 4 Pin + grounds
  - Reduced 40 pins of parallel ATA (PATA) to 4 pins to increase speed
- Speeds (Maximum)
  - PATA: 133 MB/s
  - SATA-I: 125 MB/s
  - SATA-II: 250 MB/s
  - SATA-III: 500 MB/s

Interface

- PCI Express (PCIe) v2.0
  - 500 MB/s per lane
    - Up to 16 lanes
  - Very low power and broad hardware support
- Regarding SSD
  - PCIe SSDs are used only for extremely high performance applications due to their enormous cost

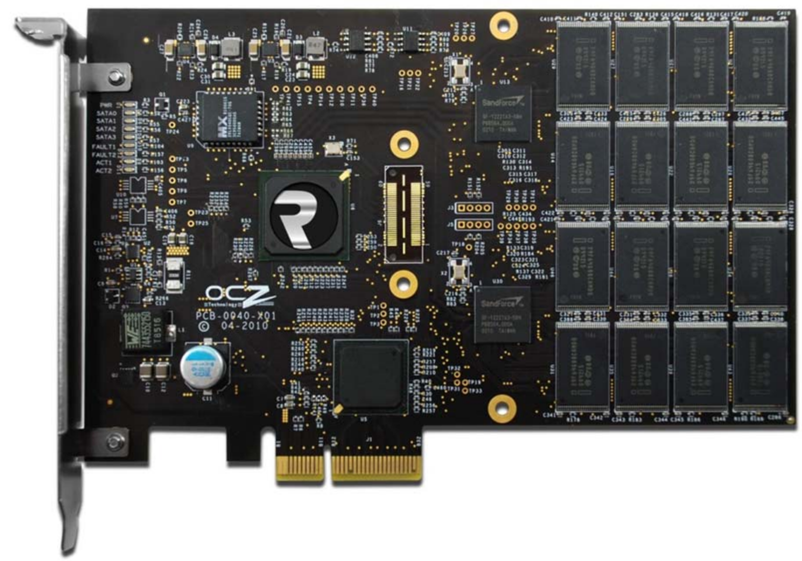

Latencies

- HDD
  - Seek time – amount of the necessary to move the head to the desired cylinder
  - Rotational time – amount of time needed for the media to rotate to the correct sector
- SDD
  - Single fetch latency that is orders of magnitude less than HDDs latency

Latency

- Video Demonstration
  - 6:40 -> 7:54

SSD Architecture

- SSDs contain a number of NAND flash components
  - 10 to upwards of 60 or 70

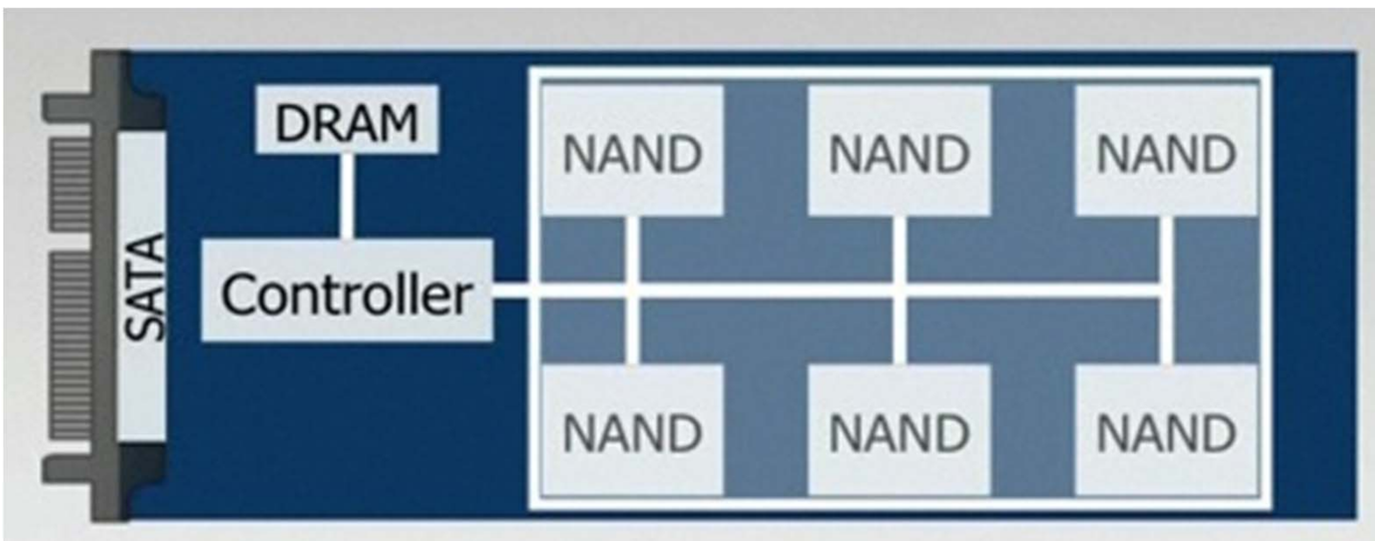

SSD Architecture

- Controller
  - Takes the raw data storage in the NAND flash and makes it look and act like hard disk drive
  - Contains the micro controller, buffer, error correction, and flash interface modules

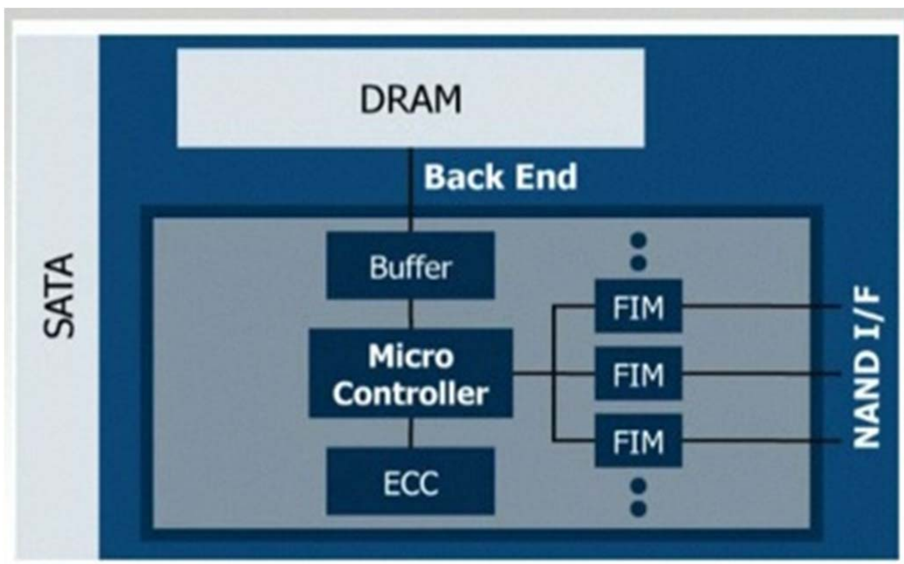

Controller Components

- Flash interface modules (FIMs) physically and logically connect the controller to the NAND flash devices
  - FIMs have the ability to communicate with multiple NAND flash devices and therefore performance can be increased by adding additional FIMs
- Micro Controller – a processor inside the controller that takes the incoming data and manipulates it
  - Stripping any errors
  - Making sure it is correctly mapped
  - Putting it into the flash or retrieving it from the flash
- DRAM Cache – Reasonable amount of very low latency memory that gives the processor some room to work

NAND Flash Media

- NAND Flash Media contains NAND cells arranged in multiple planes
  - The planes allow for parallel access to the NAND
  - They also allow for interleaving
- The data moves in and out through a cache element

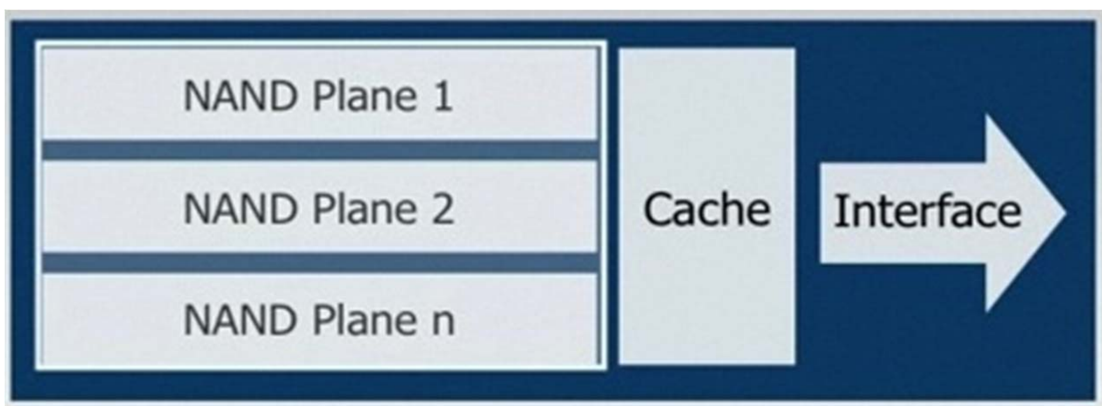

SSDs MLC vs. SLC

- MLC (Multi-Level Cell) – Cheaper yet slower and slightly less reliable
- SLC (Single Level Cell) – Faster and more reliable, yet more expensive than SLC
- Advantages of SLC are dwindling due to advancements in controller design, which mitigate the disadvantages of MLC

Advantages Summary

- Super low latency
  - Zero seek time
- Very fast read and write speeds
- Physically more robust
  - Shock resistance (1500Gs+)
  - Zero moving parts
    - Completely silent
    - Low power consumption
- Excel at small/short reads and writes
- Immune to data fragmentation

Performance Comparison

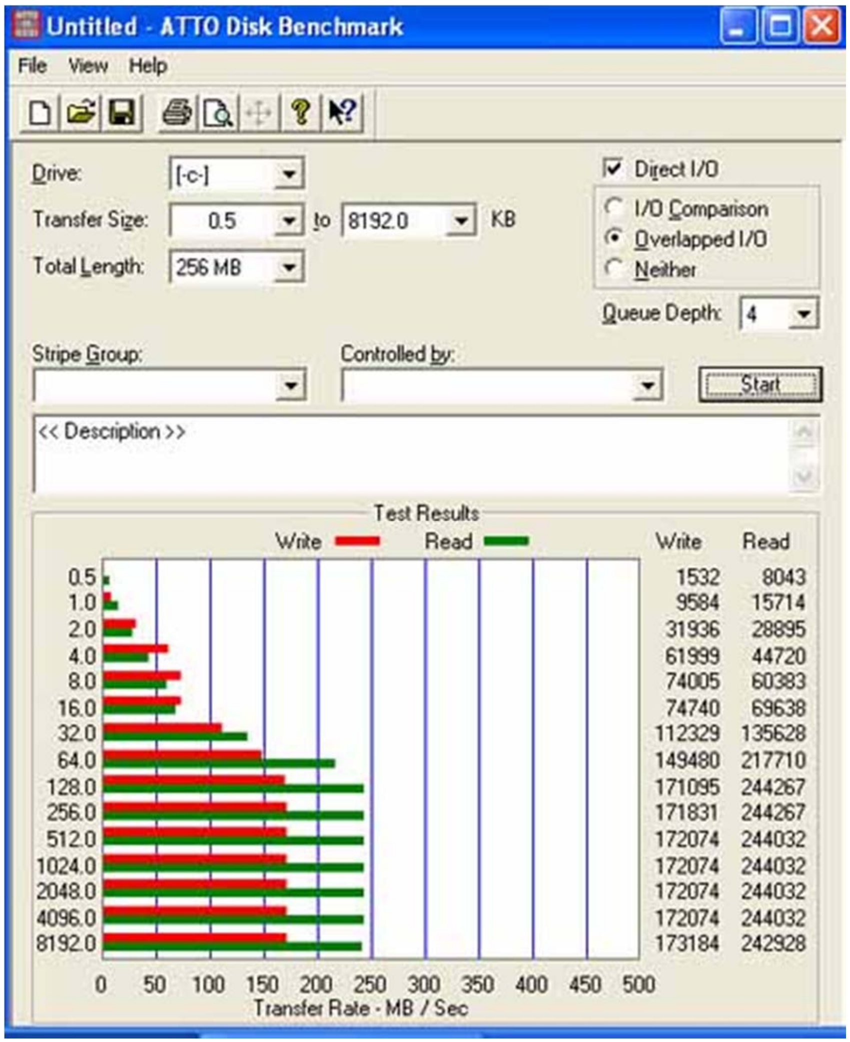

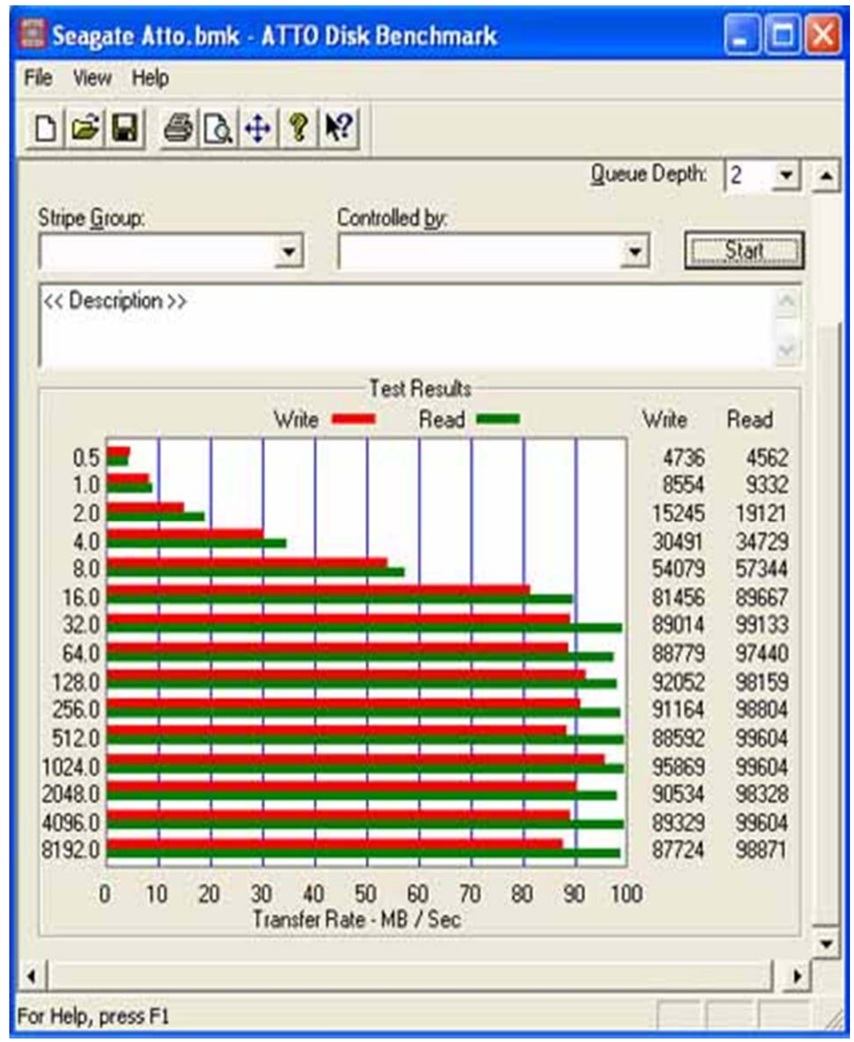

Performance Comparison

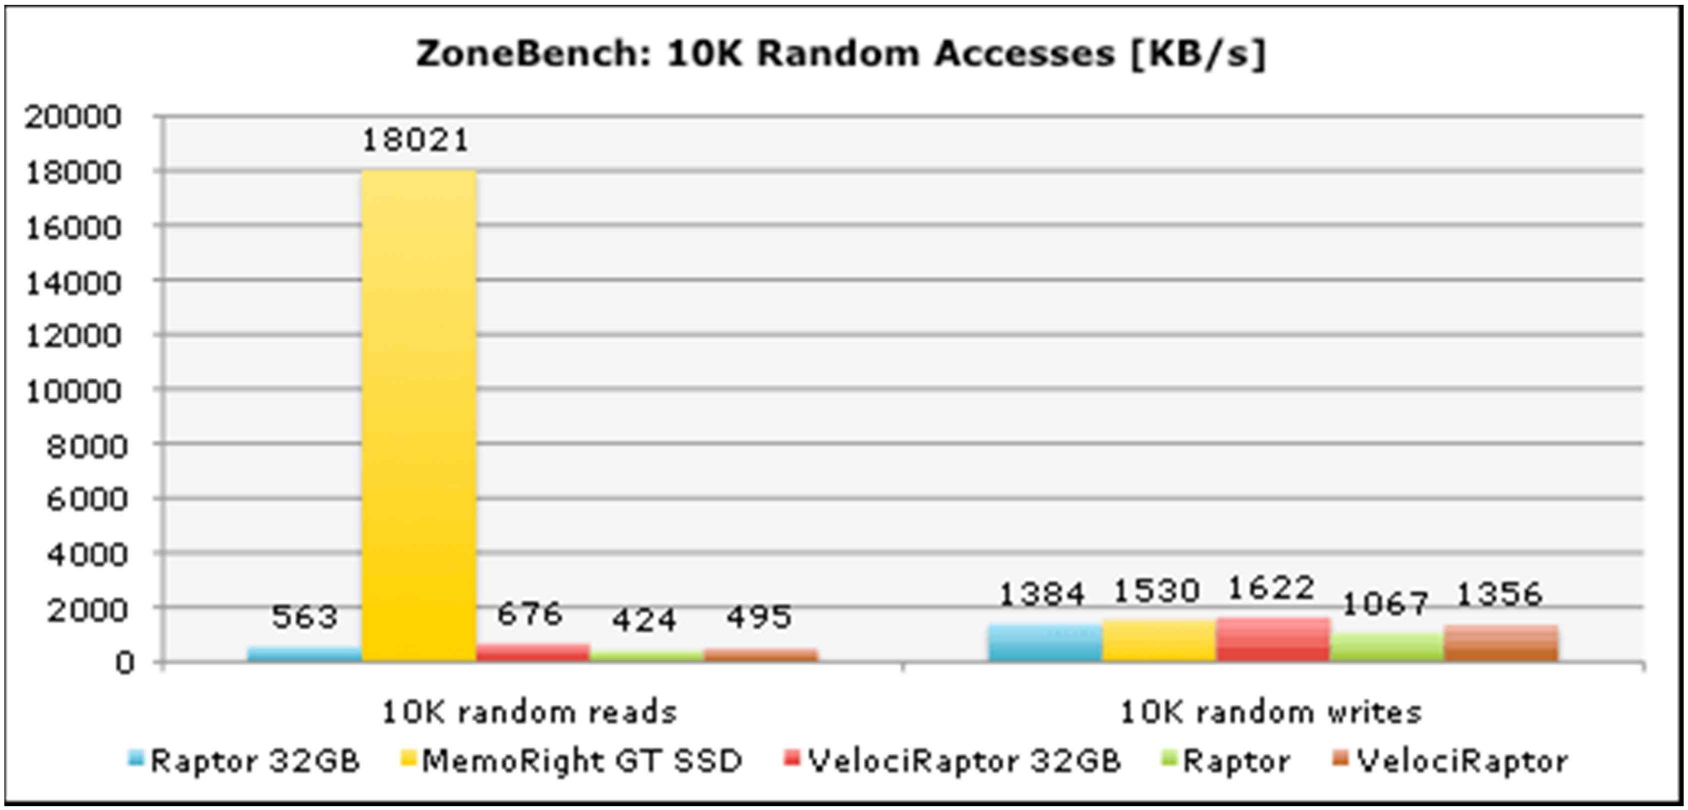

Disadvantages

- Cost per GB is much higher than HDDs
  - 64GB SSD $125-$150
  - 2TB HDD $90-100
- General Size
  - 3TB consumer 3.5’’ HDDs are available (relatively common)
  - 1TB consumer 3.5’’ SSDs are available (very rare and expensive)
- Limited write cycles
  - 1 to 2 million write cycles before wear out for MLC
  - Up to 5 million write cycles before wear out for SLC

Conclusions

- Although cost prohibitive, for performance applications SSDs hold a great advantage over platter drives
- Common consumer setups take advantage of SSDs for program files while using larger cheaper platter drives to store media and other general storage
  - Hybrid drives
- SSDs represent the evolution toward alleviating the bottleneck that is data storage in present day systems

Questions?