SSDs
A Shift in Data Storage

Todd Dinkelman
Micron Technology, Inc.

Agenda

- SSD Defined
- Quick Comparison
- Why MLC
- Reliability/Endurance
- Summary

What is Being Called an SSD?

- USB interface embedded solution
- Module-based
- IDE interface card w/o enclosure
- 1.8-inch and 2.5-inch, 32GB and 64GB SSD for notebook and performance computing
- Module/card with custom form factor, density, and interface

The Complexity of HDDs

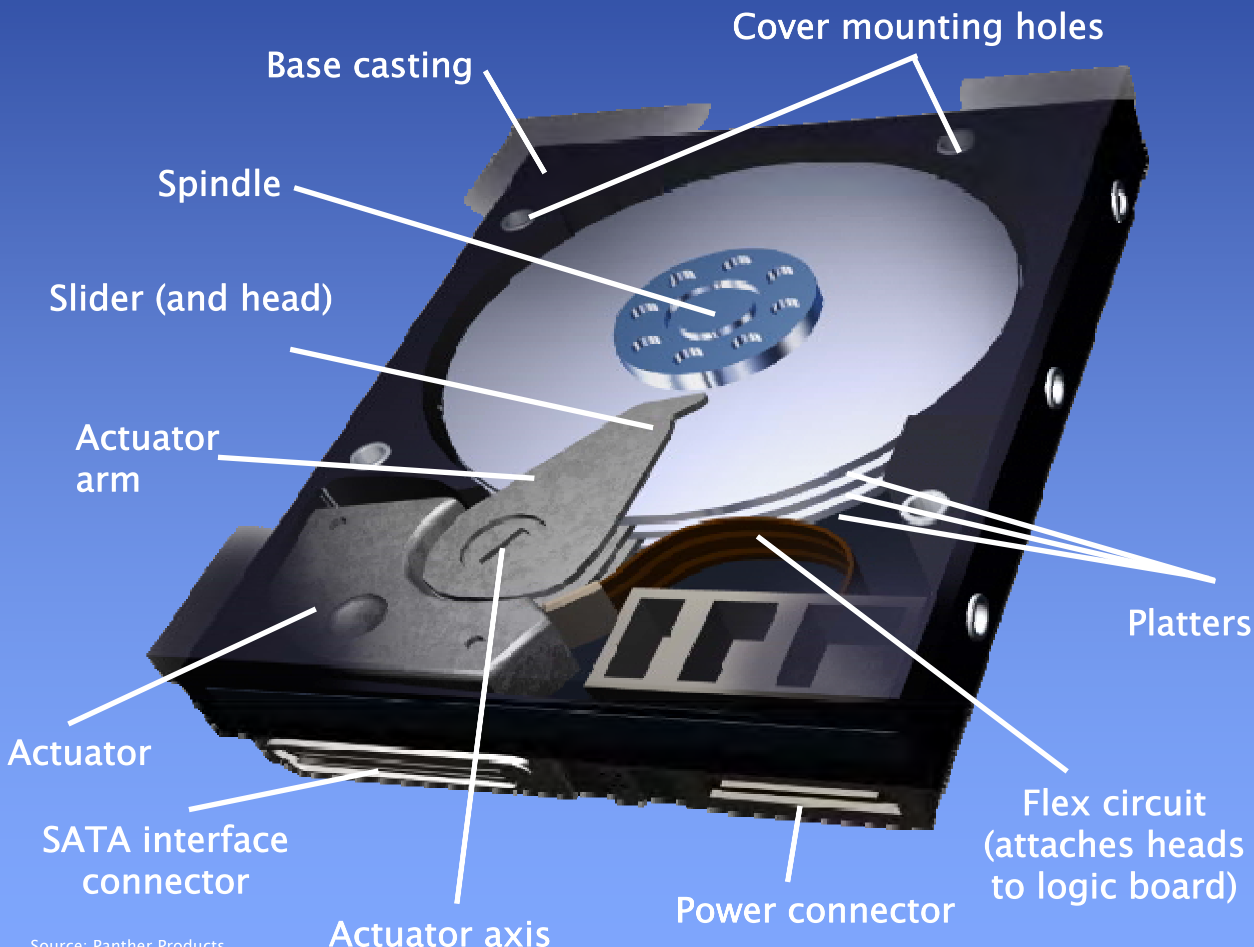

Source: Panther Products

- HDD Advantages
  - Density
  - Price/GB

The Simplicity of SSDs

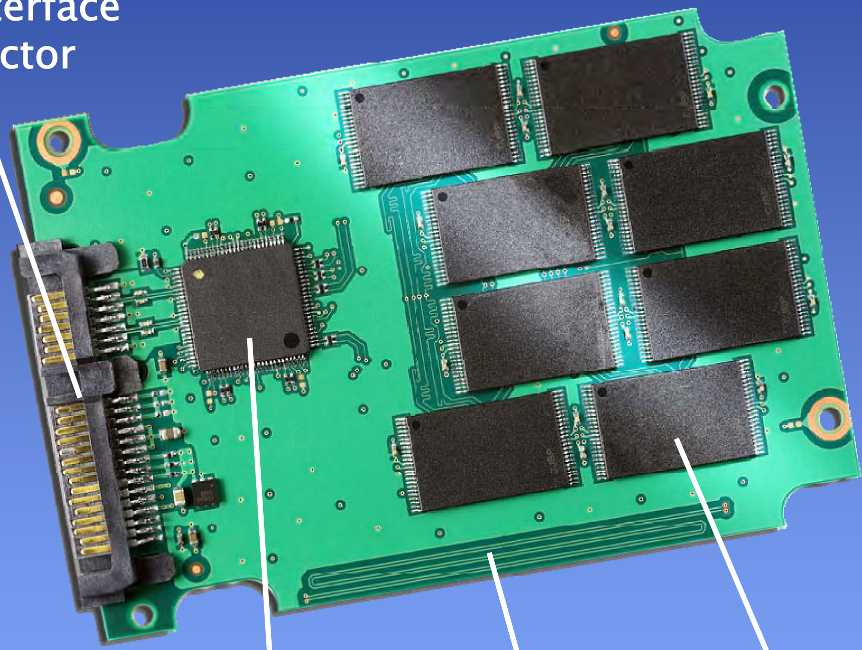

- SSD Advantages
  - Performance
  - Size
  - Weight
  - Ruggedness
  - Temperature Range
  - Power

SSDs vs. HDDs

	SSDs	HDDs
Capacity		✓
Performance	✓	
Reliability	✓	
Endurance	✓	✓
Power	✓	
Size	✓	
Weight	✓	
Shock	✓	
Temperature	✓	
Cost per bit		✓

- **Due to recent advances in NAND lithography, SSD densities have reached capacities for mass market appeal**
- **SSDs offer many features that lead to improved user experiences**
- **Early shortcomings regarding reliability and endurance are being overcome**

SSDs in Computing

CPU

Relative Latency		Relative Cost/Bit
1	L1 Cache	1,800
2.5	L2 Cache	1,400
1,200	DRAM	10
25,000	SSD	3
25,000,000	HDD	1

NAND

NAND Flash closes the latency gap

Cost/bit data as of November 2007

NAND in Computer Architecture

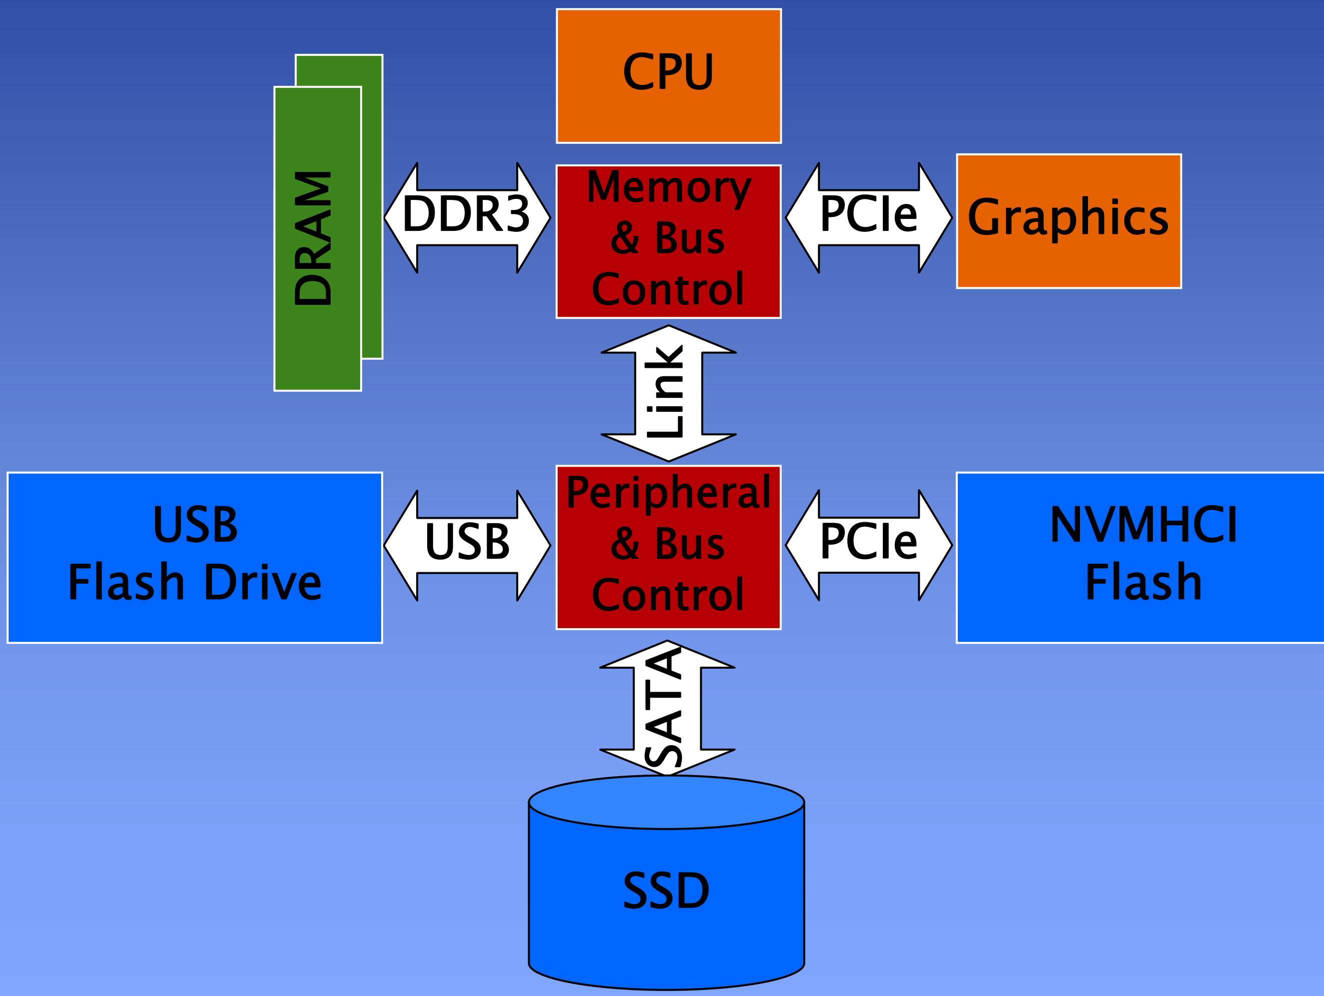

MLC and SLC Differences

- SLC
  - Single-level cell
    - One bit per cell
- MLC
  - Multi-level cell
    - 2 bits per cell – today
    - 3 and 4 bits per cell – future
- Endurance
  - SLC is typically 10 times better than MLC
- Performance
  - SLC provides ~2X the write performance of MLC
- Price
  - SLC-based products have better than 2X the price/GB compared to MLC

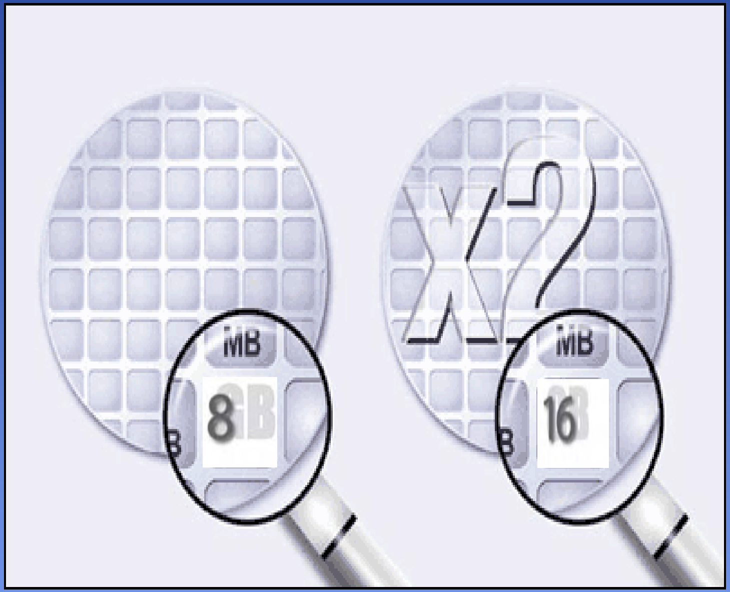

SSD Market Trends

- Improvements in controller technology
  - Moving from CompactFlash architectures to true SSD controllers
- Notebooks and PCs
  - Migration to MLC
    - Light usage model
    - Cost, size, and performance are all important
  - Value Proposition
    - Better than desktop performance in an ultra-light notebook

SSD Average Price/GB

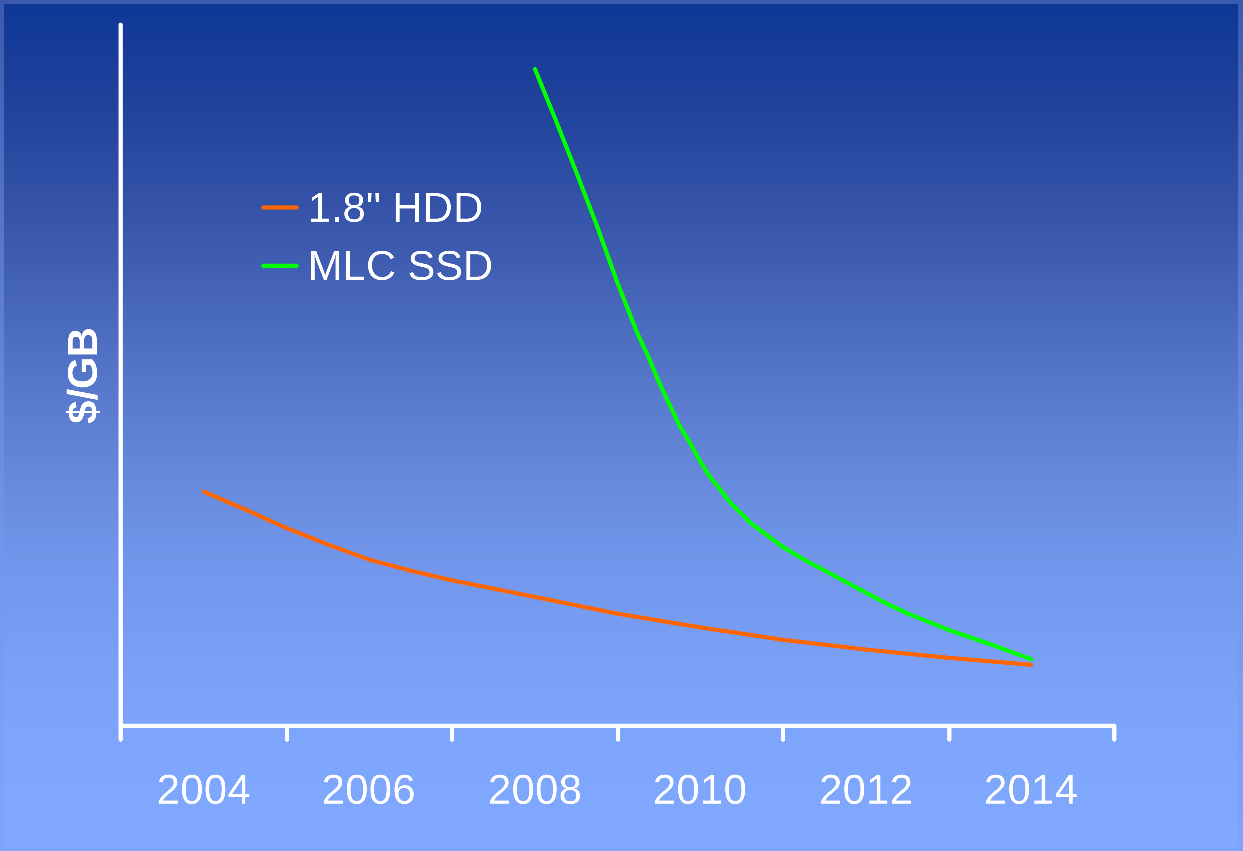

NAND Reliability and Endurance

	Effect	Description	Observed As	Management
Reliability	Program Disturb	Cells not being programmed receive charge via elevated voltage stress	Increased read errors immediately after programming	ECC and Block Management
	Read Disturb	Cells not being read receive charge via elevated voltage stress	Increased read error at high number of reads	ECC and Block Management
	Data Retention	Charge loss over time	Increased read errors with time	ECC and Block Management
Endurance	Endurance/ Cycling	Cycles cause charge trapped in dielectric	Failed program/erase status	Retire Block

NAND Error Rate

- Bit Error Rate
  - Failing bits corrected with appropriate levels of ECC
  - Correctable bit errors do not result in data loss
- Raw Bit Error Rate (RBER): Bit error rate prior to ECC
- Uncorrectable Bit Error Rate (UBER): Bit error rate after ECC
- UBER is projected using the measured RBER and specific level of ECC

SSD Reliability and Endurance

- SSD reliability has two parts
  - MTBF
    - Measure of time between failures due to manufacturing or component defects
    - 2 million hours is typical for Micron SSDs
  - Endurance
    - SSDs all wear out due to data writes
    - Indication of drive life based on a usage condition
    - Micron SSDs are specified to last 5 years under predefined usage conditions
    - Usage conditions vary for consumer and performance products

Endurance Factors

- Wear-leveling efficiency
- Write amplification
- NAND cycles
- SSD densities

Wear-Leveling

- Dynamic wear-leveling uses the available free space and the incoming data to equally wear each of the physical blocks
  - Static data is not included in the available pool of wear-leveling blocks, leaving a portion of the drive with no wear
- Static wear-leveling considers all physical blocks in the SSD, regardless of content, and maintains an even level of wear across the entire drive

Wear-Leveling Example

- MLC devices can typically support 10,000 cycles per block
- If you erased and reprogrammed one block every second, you would exceed the 10,000 cycling limit in just 3 hours!

  60 x 60 x 3 = 10,080

- Rather than cycling the same block, wear-leveling involves distributing the number of blocks that are cycled

Wear-Leveling Example *(continued)*

- An 8GB MLC-based SSD contains 32,768 independent blocks (each block is 256KB of data)
- If we took the previous example and distributed the cycles over all 32,768 blocks, each block would have been programmed once after 9 hours
- If you provided perfect wear-leveling on an 8GB drive, you could erase and program a block every second, every day for over 10 years!

$$\frac{10{,}000 \times 32{,}768}{60 \times 60 \times 24} = \frac{327{,}680{,}000}{86{,}400} = 3{,}792 \text{ days} = 10.38 \text{ years}$$

Write Amplification

- Additional write overhead due to garbage collection and wear-leveling
  - Write Amplification = Flash Writes / Host Writes
- Influence on host data patterns
  - Large number of small random transfers increases write amplification
- Influenced by number of spare blocks and static data
  - Keeping a percentage of the drive as spares will bound the write amplification

Typical PC File Transfers

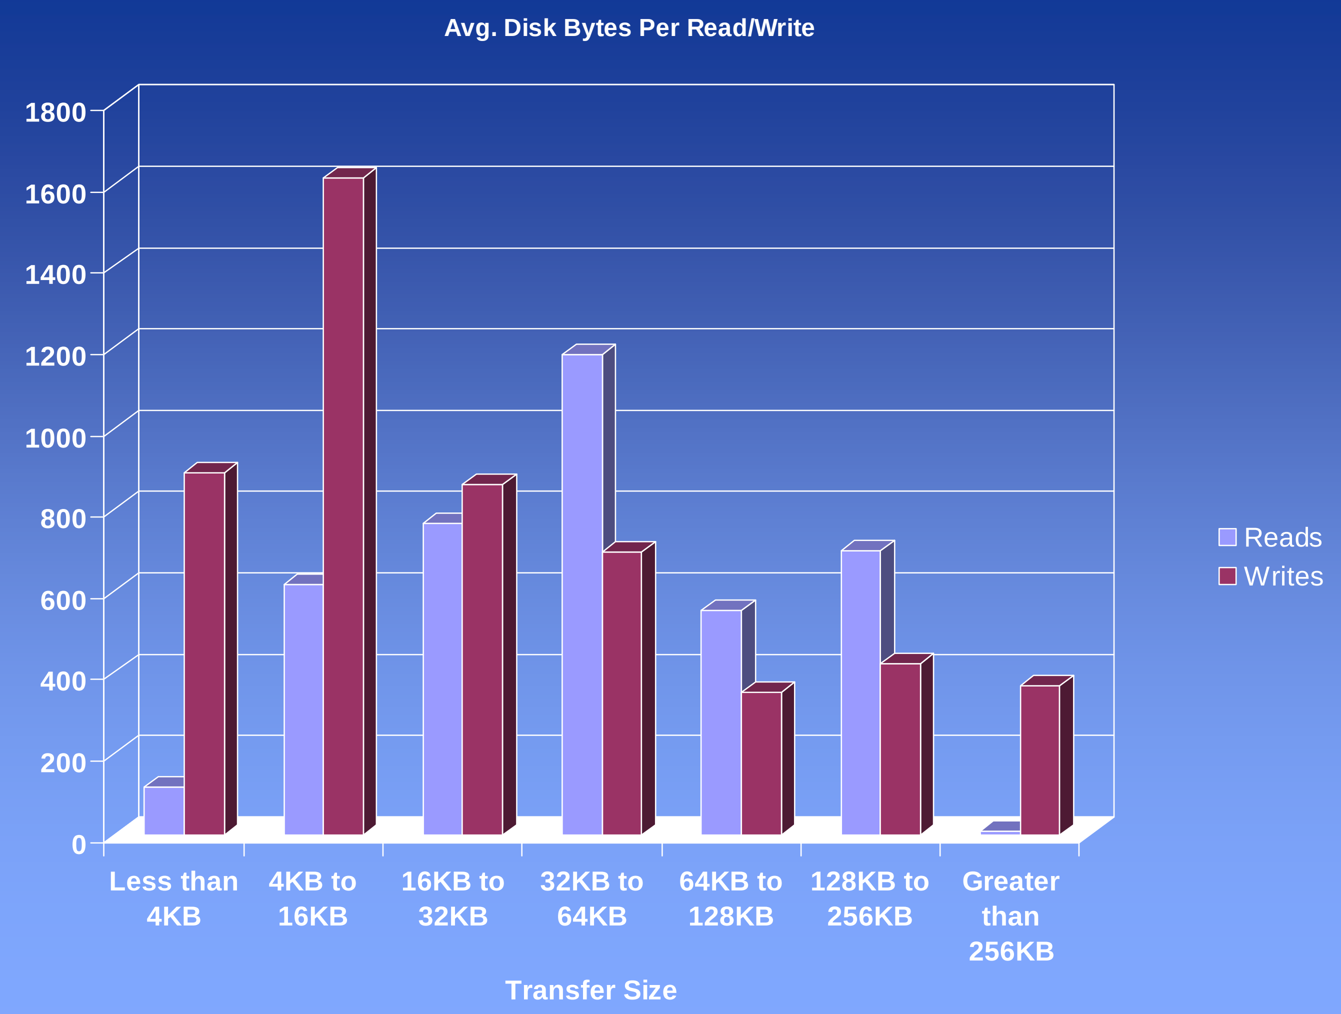

SSD Endurance Calculation

- Write amplification and wear-leveling efficiency must be accounted for when calculating SSD lifetime

$$\text{Life in Years} = \frac{\text{NAND Cycles} * \text{SSD Capacity}}{\text{Amplification Factor} * \text{GB/Year}}$$

SSD Lifetime

- Given a <u>capacity point</u>, <u>NAND endurance</u>, and <u>writes-per-day</u>, the upper limit on SSD lifetime is defined
  - Example: A 64GB SSD, 10K PROGRAM/ERASE cycles and write duty of 350 GB/day will wear out in 5 years
- Uneven wear within a NAND Flash device may cause 'hot spots'
  - Wear-leveling of NAND is necessary
- Improper NAND data management may cause lifetime to be significantly less than the limit
  - A Flash ERASE operation must occur before a page can be rewritten
  - To improve overall throughput, erase granularity is much larger than write granularity
  - LBA traffic pattern can cause lots of extra data movement

SSD Endurance

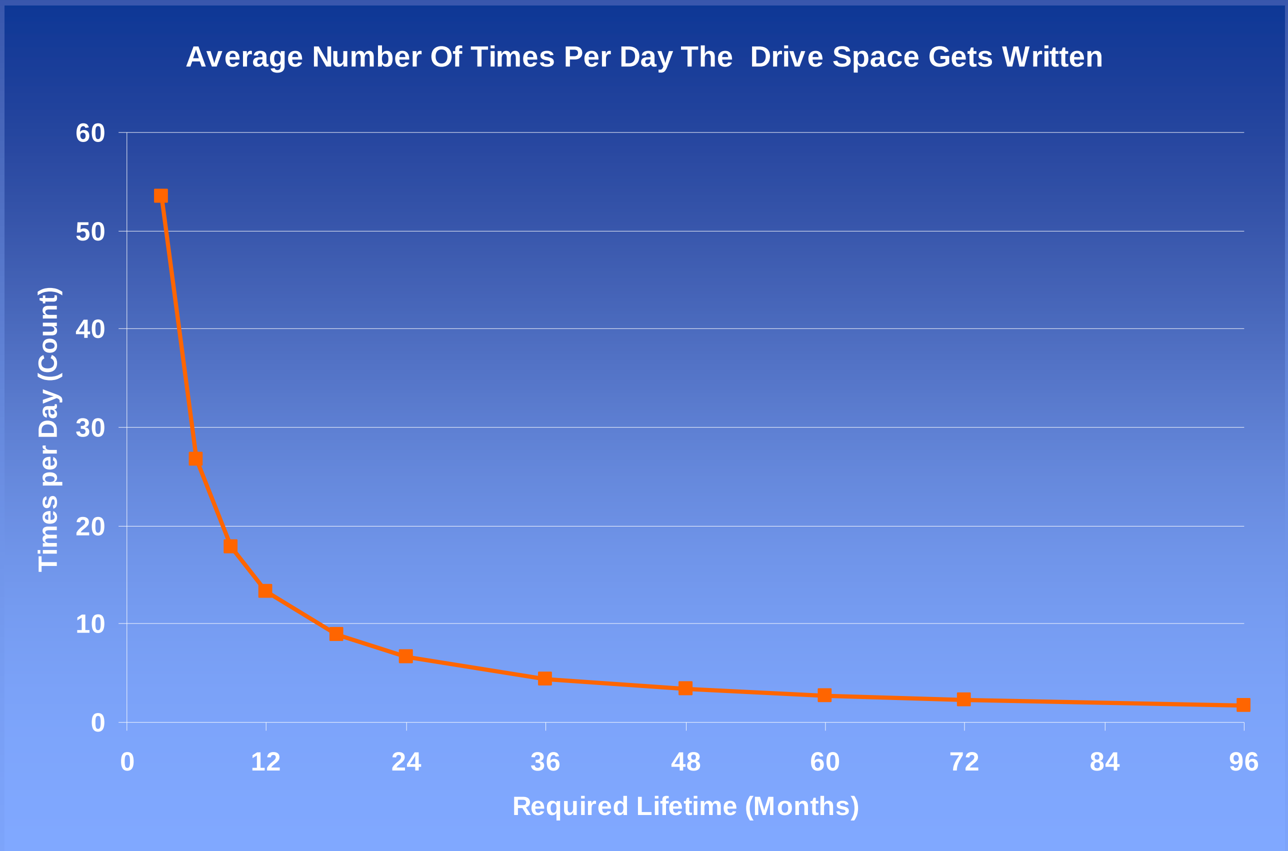

33% WRITEs
2K random IOPs
64GB drive
(55GB user-accessible
20% reserved spares

SSD Capacity Effect on Endurance

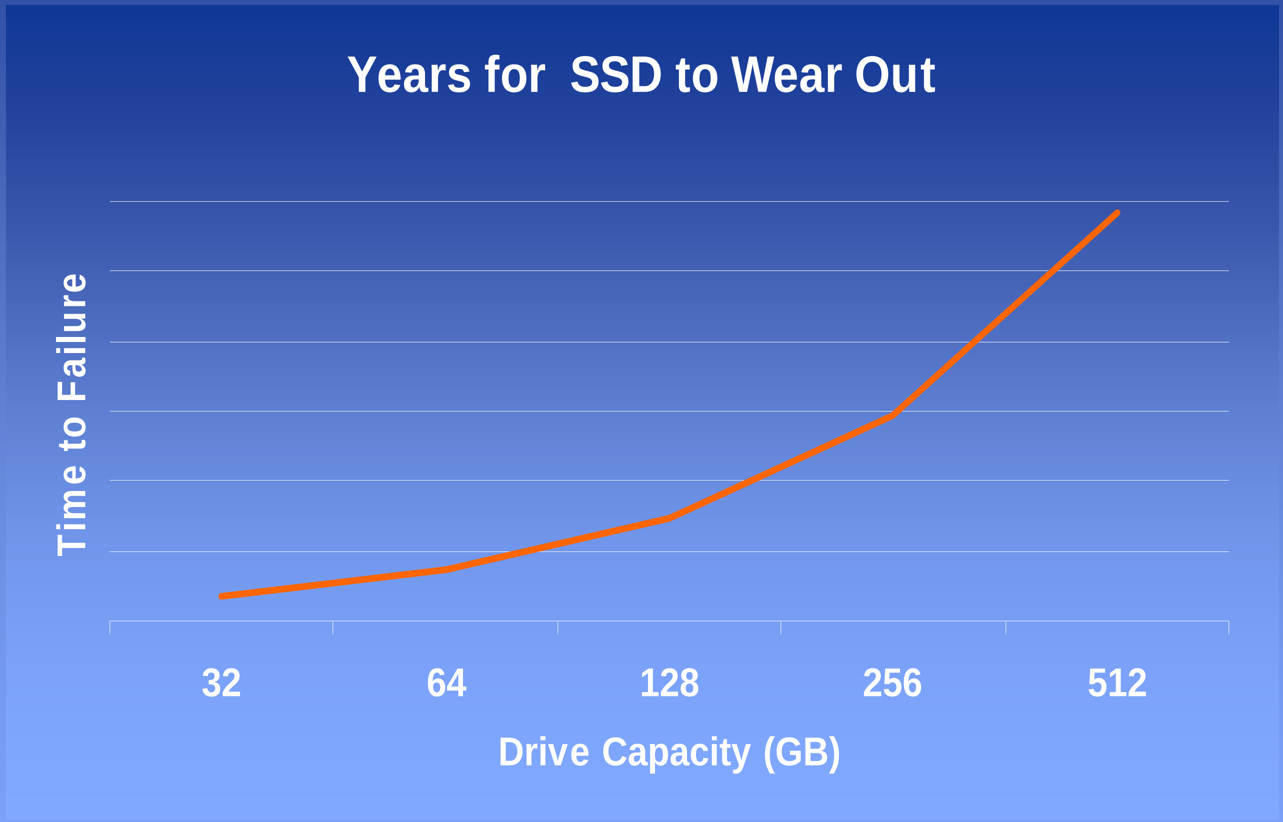

- SSD life will double with every doubling of capacity

Summary

- SSDs bring many advantages to storage
  - Performance
  - Power
  - Reliability
  - Environmental ruggedness
- Different products for different markets
  - MLC-based SSDs for PCs
- Reliability has two aspects
  - Endurance and MTBF

Thank You