Memoright Corporation

The Age of Application Specific SSD

By

Eric Kao

IS IT AN OVER-STATEMENT?

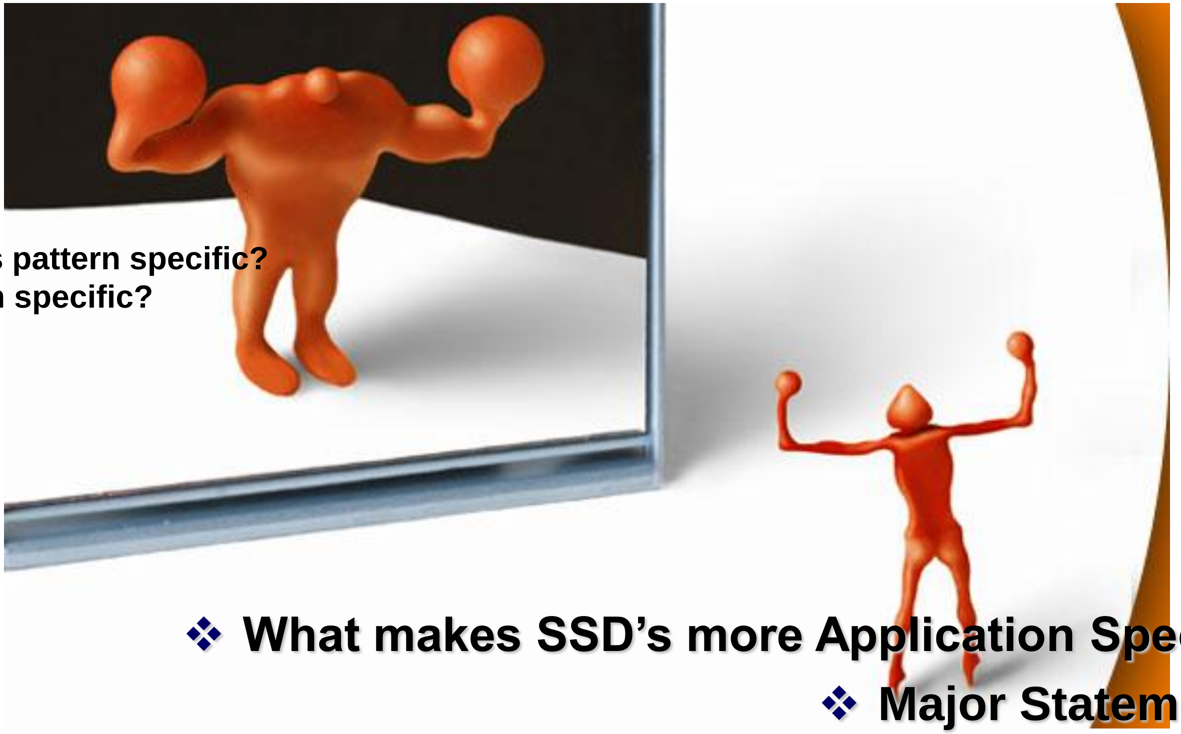

- Access pattern specific?
- System specific?

❖ What makes SSD's more Application Specific than HDD's?

❖ Major Statements for This Talk

❖ Firmware defines the personality - 4 case studies

❖ How Does the Future Look Like?

What makes SSD’s more A.S. than HDD’s?

- **These stuffs matter... but Firmware is the key subject today.**

Interface:
PATA, SAS & SATA, FC

Physical Design:

Interface:
PATA, SAS, SATA, FC, mSATA, mPCIE, PATA ZIF, CF, eMMC…etc.

Physical Design: (unlimited)

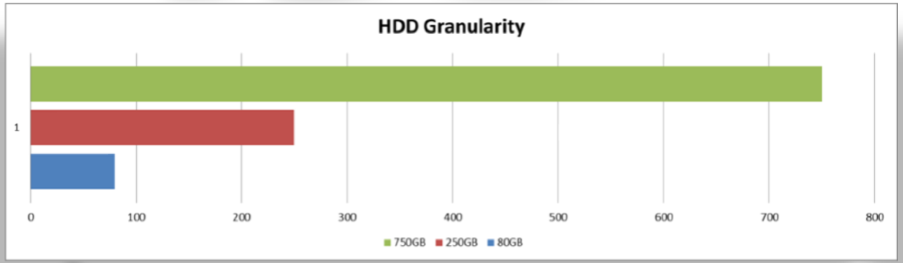

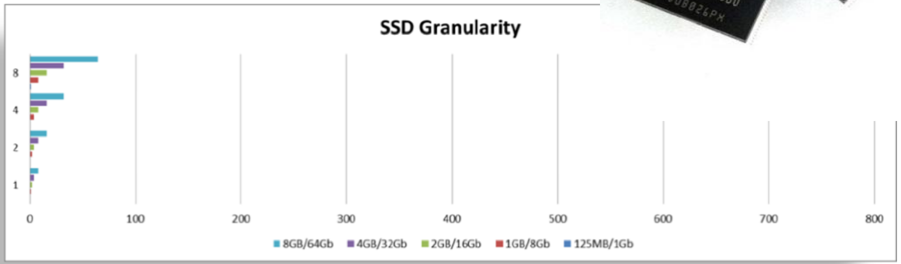

Major Statements for this talk

- What sets SSD's personality apart is the firmware design.
- Once a drive is installed, it stays in that application for its entire life.
- Therefore, it's worthwhile to create different FW stacks optimized for certain access patterns.
- Furthermore, it is not only worthwhile but also necessary because it is impossible to have one FW design which can satisfy all criteria from all applications.

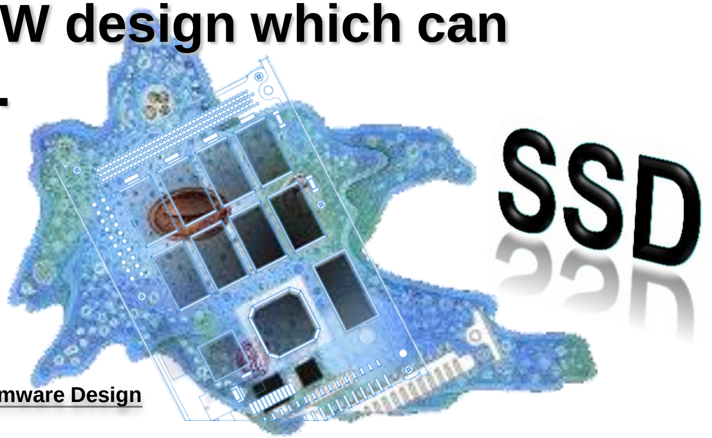

Firmware Design

Is it an over-statement?

HDD's have diversified….

Application	HDD	Special Design
Enterprise	WD VelociRaptor	High RPM (7200 ~ 15K RPM) for high performance
	WD RE SAS	Lower RPM (6400RPM) for better reliability
	WD RE4/RE4-GP	High speed interface (SAS/Fiber Channel)
	Seagate SAVVIO	Better ECC data correction
	Seagate Constellation/ES	Low power consumption design
	Hitachi Ultrastar Series	Anti-vibration mechanical design
Surveillance	Seagate Pipeline HD	Longer MTBF design
	Seagate SV35 Series	Wide temp design and material
	WD AV/AV-GP/AV-25	Firmware optimized for video streaming
	Hitachi Cinemaster Series	Lower RPM (4200~5400 RPM) for better reliability
Automotive	Hitachi Endurastar Series	High capacity disk
	Seagate EE25 Series	Wide temp design and material (-30°C to +85°C)
		Anti-vibration mechanical design (shock sensor)
IPC	Seagate SAVVIO Series	Anti-humidity design
	Seagate Cheetah	ECC/Self encryption, improving security and system performance
	Hitachi Endurastar Series	Wide temp design and material (-30°C to +85°C)
		Anti-vibration mechanical design (shock sensor)
Consumer/DVR/Media Center	WD AV/AV-GP/AV-25	Anti-humidity design
	Hitachi Cinemastar Series	Quiet operation design, low noise motor and mechanical design
	Seagate Pipeline HD	High capacity disk
	Seagate SV35 Series	Firmware optimized for video streaming
		Larger cache memory
PC	WD Caviar Blue/Green/Black	Low power consumption design
	Seagate Barracuda Series	ECC/Self encryption, improving security and system performance
	Hitachi Deskstar Series	Larger cache memory/SSD cache
		7200RPM performance
NB	WD Scorpio Blue/Black	4K format to optimize system performance
	Seagate Momentus Series	Ultra slim form factor mechanical design (7mm low-profile)
	Hitachi Travelstar Series	Shock-and vibration-proof ruggedness
		7200RPM performance

What makes SSD's more A.S. than HDD's?

HDD & SSD Latency comparison

HDD Latency = Rotational Latency + Track Seek Time

From 0.x ms to xx ms

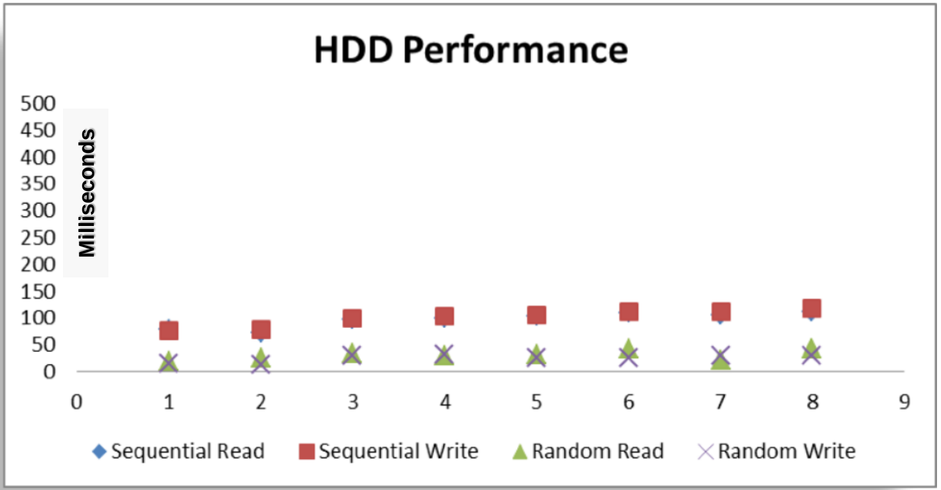

SSD Latency = **Function(** #Channels,
#CE,
NAND Flash Access Time,
NAND Flash Page Size,
Writing Strategy,
Mapping Scheme,
Garbage Collection,
Wear-Leveling,
ECC,
…)

From xx us to x second

What makes SSD's more A.S. than HDD's?

Criteria trade-off in FW Design

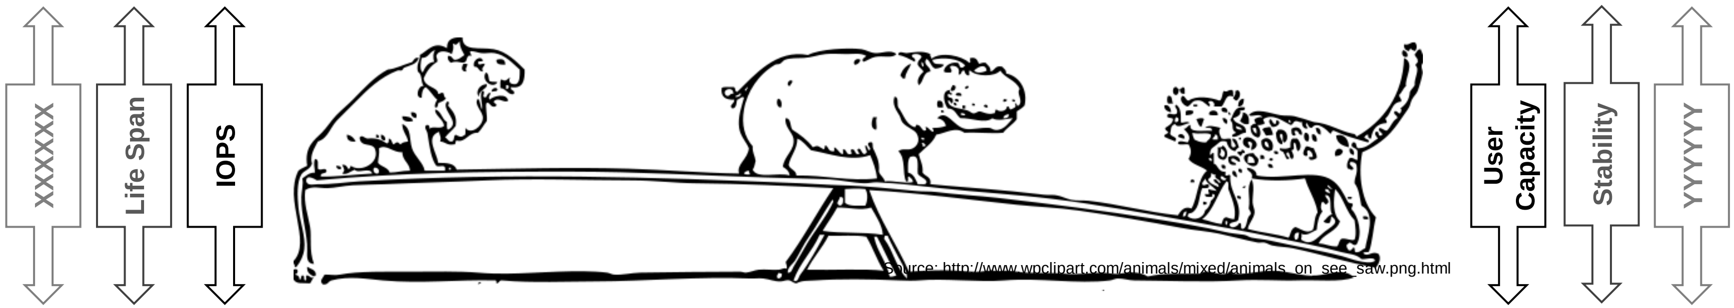

- High throughput
- Low Redundancy
- High User Capacity
- Good Performance Stability
- High Data integrity
- Long Life Span
- Manageable Latency

What makes SSD's more A.S. than HDD's?

- **SSD is actually a large array of drives with 100's of flash chips.**

HDD

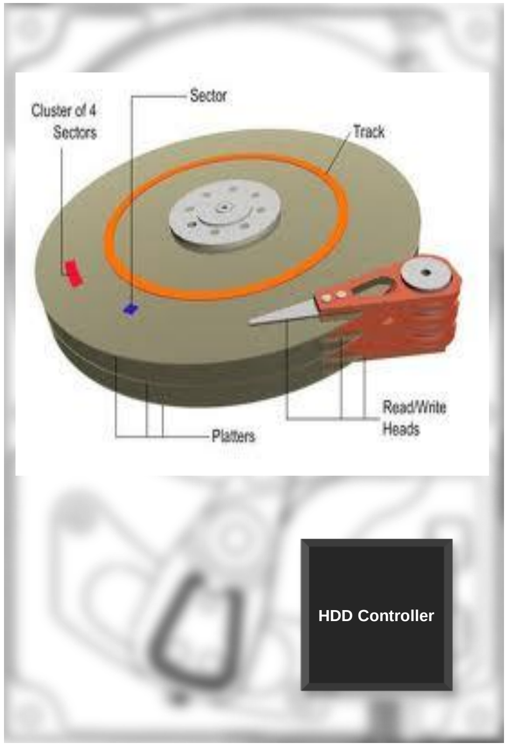

SSD

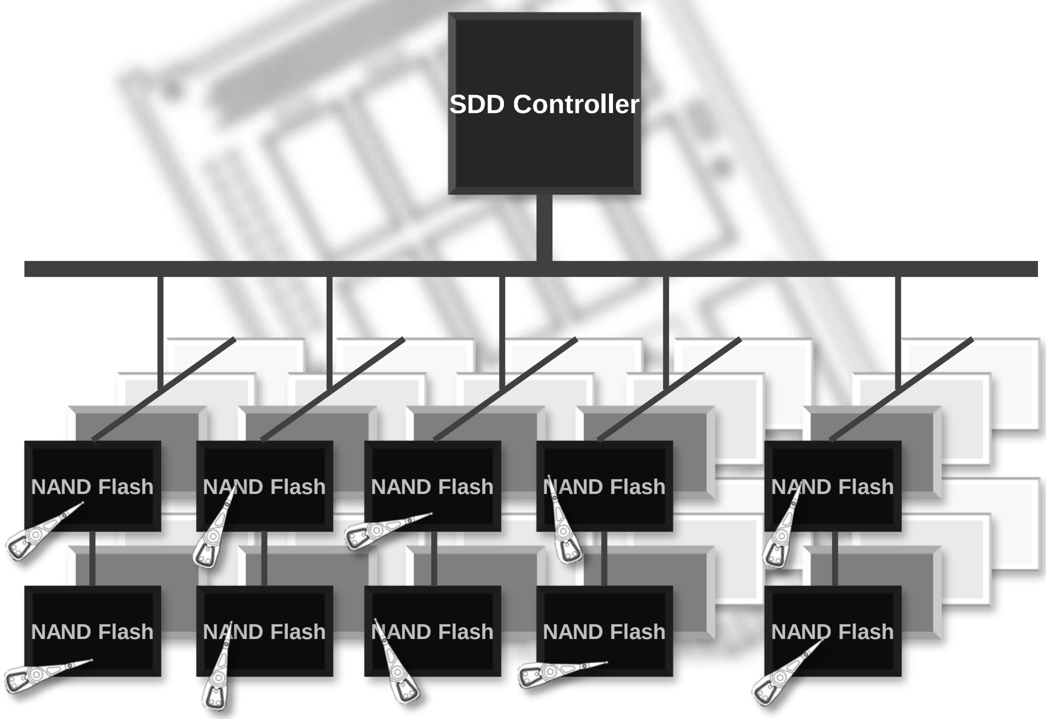

What makes SSD's more A.S. than HDD's?

- **There is conflict between flash nature and SSD design**
  - Flash blocks like to be written in full due to its block-erasure nature,

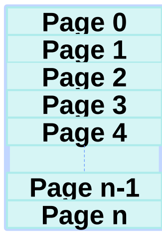

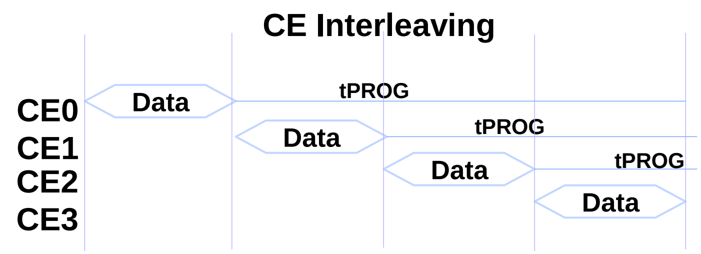

  - But interleaving scheme likes to write across as many blocks as possible, thus creates fragmentations.

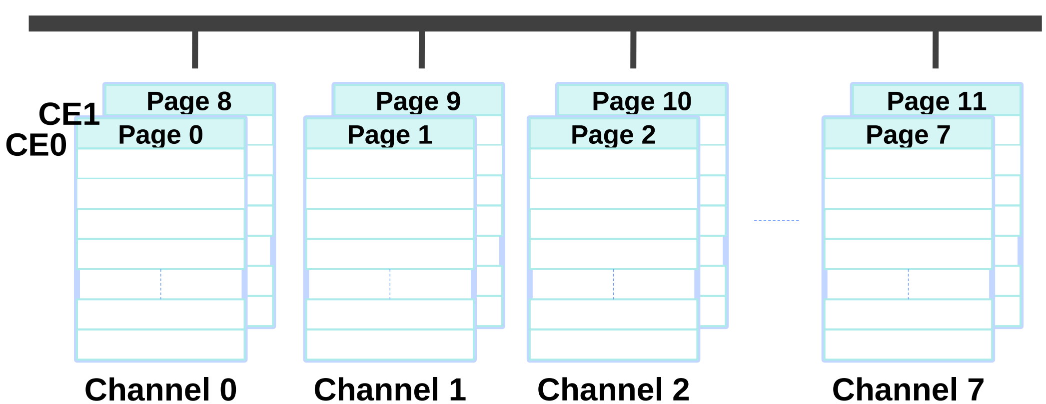

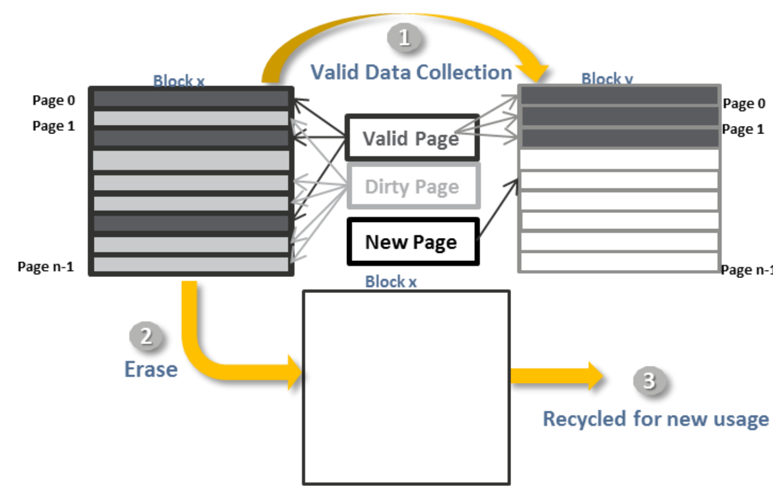

Firmware defines the personality

Case Study 1 – Digital Video Recording / Editing

- Video Stream: 40MB/s ~ 500MB/s, payload: 1MB ~ 8MB
- Audio Stream: 284KB/s ~ 1.5MB/s, payload: 8KB ~ 64KB
- Metadata: 128KB
- Array of n SSD's

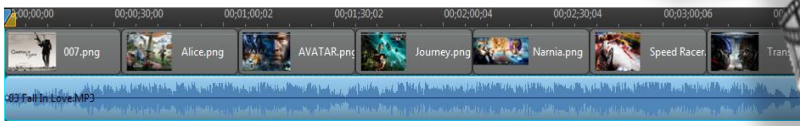

Best Choice of approach

- Block mapping, vertical writing
- Data buffering for small payload
  - Reduce fragmentation
  - Reduce write-amplification
  - Constant speed

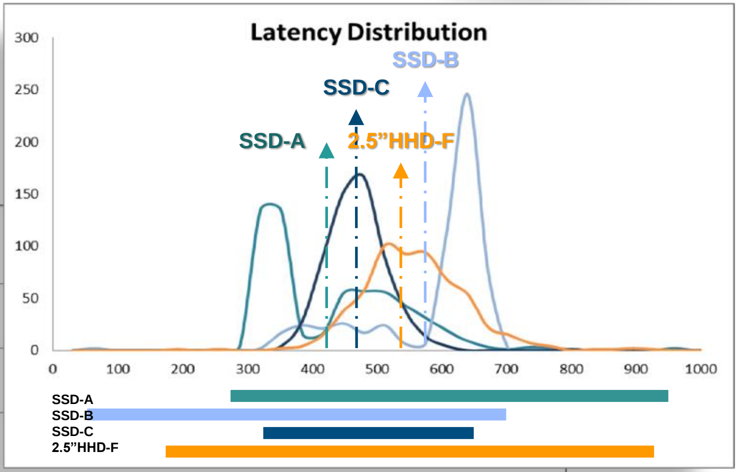

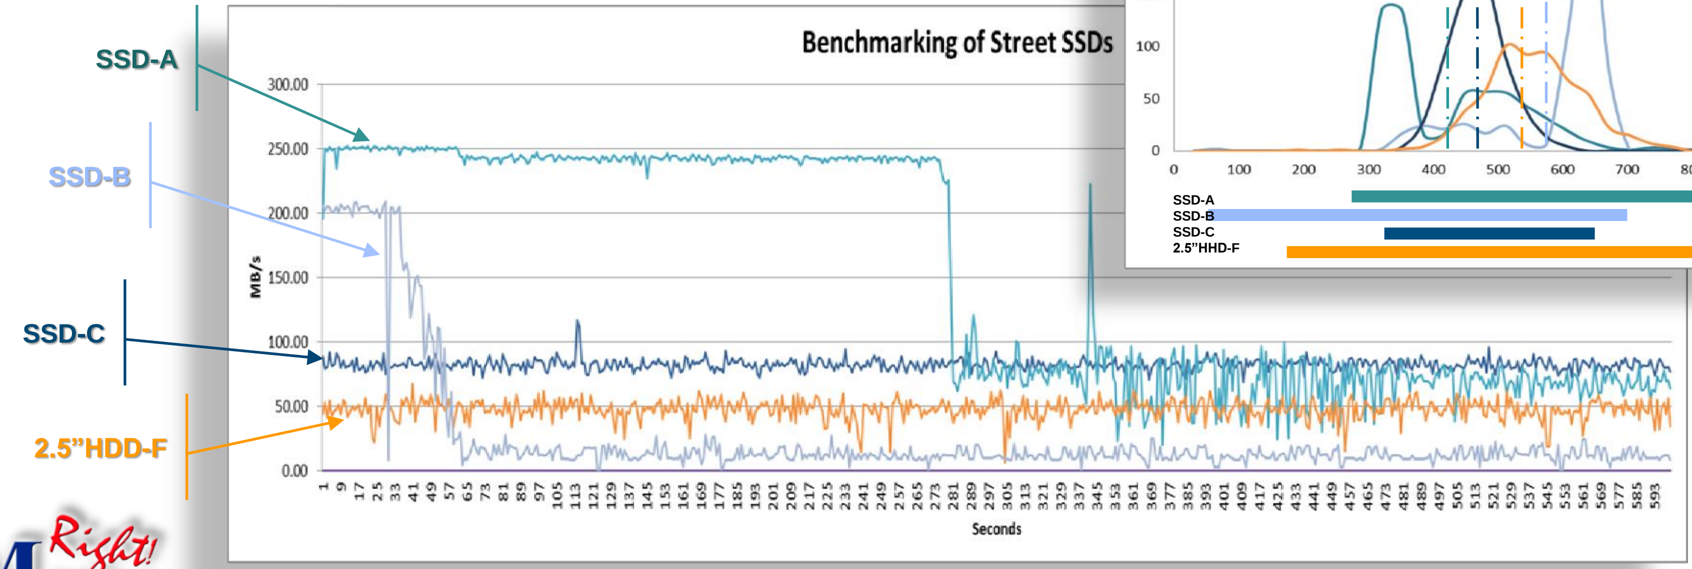

Firmware defines the personality

Case Study 2: Main Stream PC

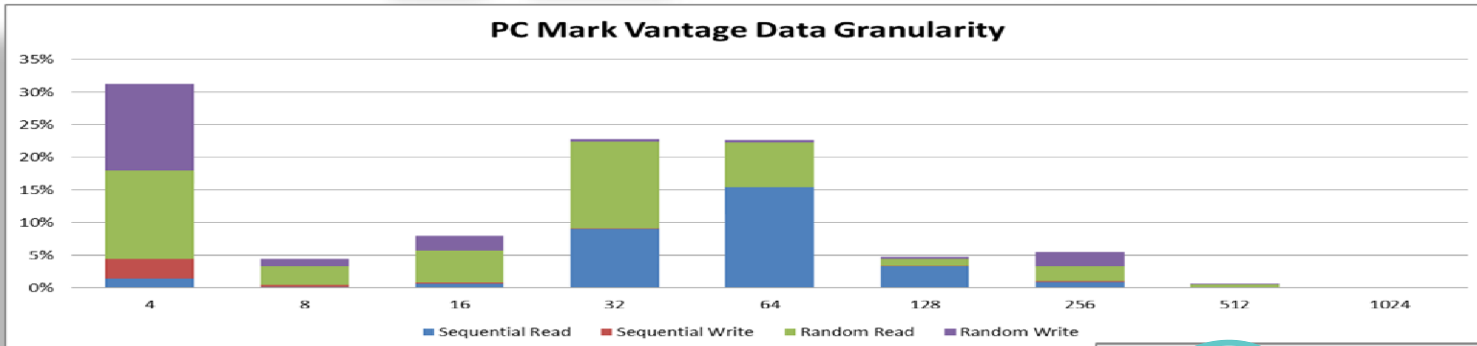

Best Choice of approach

- Pure page mapping for low WA & high IOPS
- 4KB mapping granularity
- Low over-provisioning for more user space
- Will have wide-spread latency, it's OK though.

Firmware defines the personality

- **Case Study 3 – RAID Disk Array**

Source: http://h41131.www4.hp.com/

- **Demands:**
  - Random 4KB~256KB IOPS
  - Narrow-spread latency
  - High reliability
  - High endurance
- **Best Choice Of The Approaches**
  - Page mapping for good random IOPS & low WA
  - Finer space partitioning for narrow latency
  - In-drive redundancy for better BER and die-failure recovery
  - SLC/eMLC

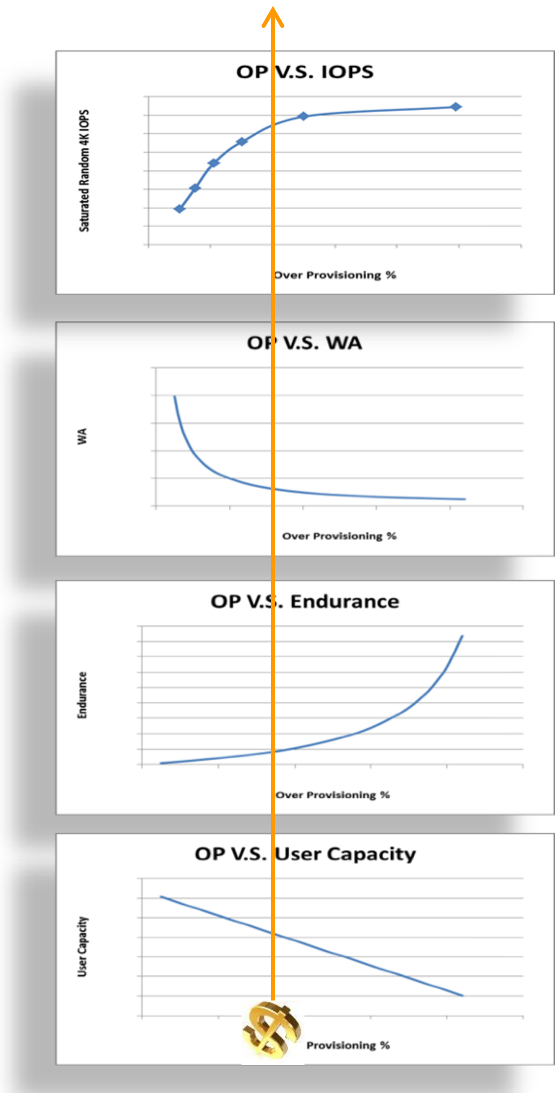

Firmware defines the personality

Case Study 4 – Linear Drive

Source: http://tape-drive-recall.com/

Demand

- Pure sequential writing by Super Blocks
- Random page read
- High MTBF

Best Choice Of The Approaches

- Super block mapping with bad blocks management
- Chip-level redundancy

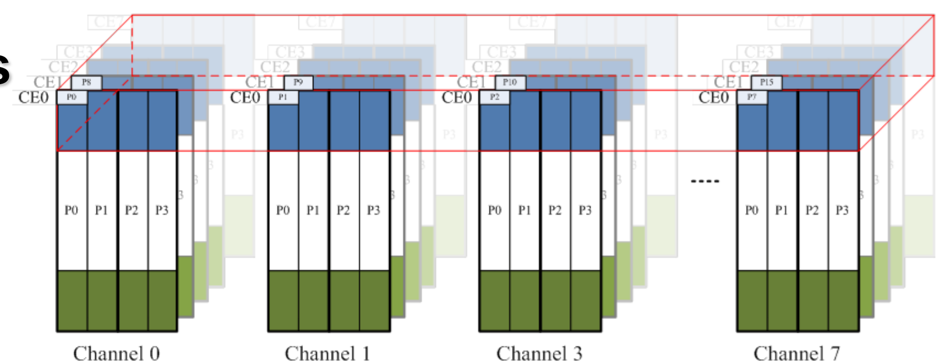

How does the future look like?

Case Study 5: Tablet PC

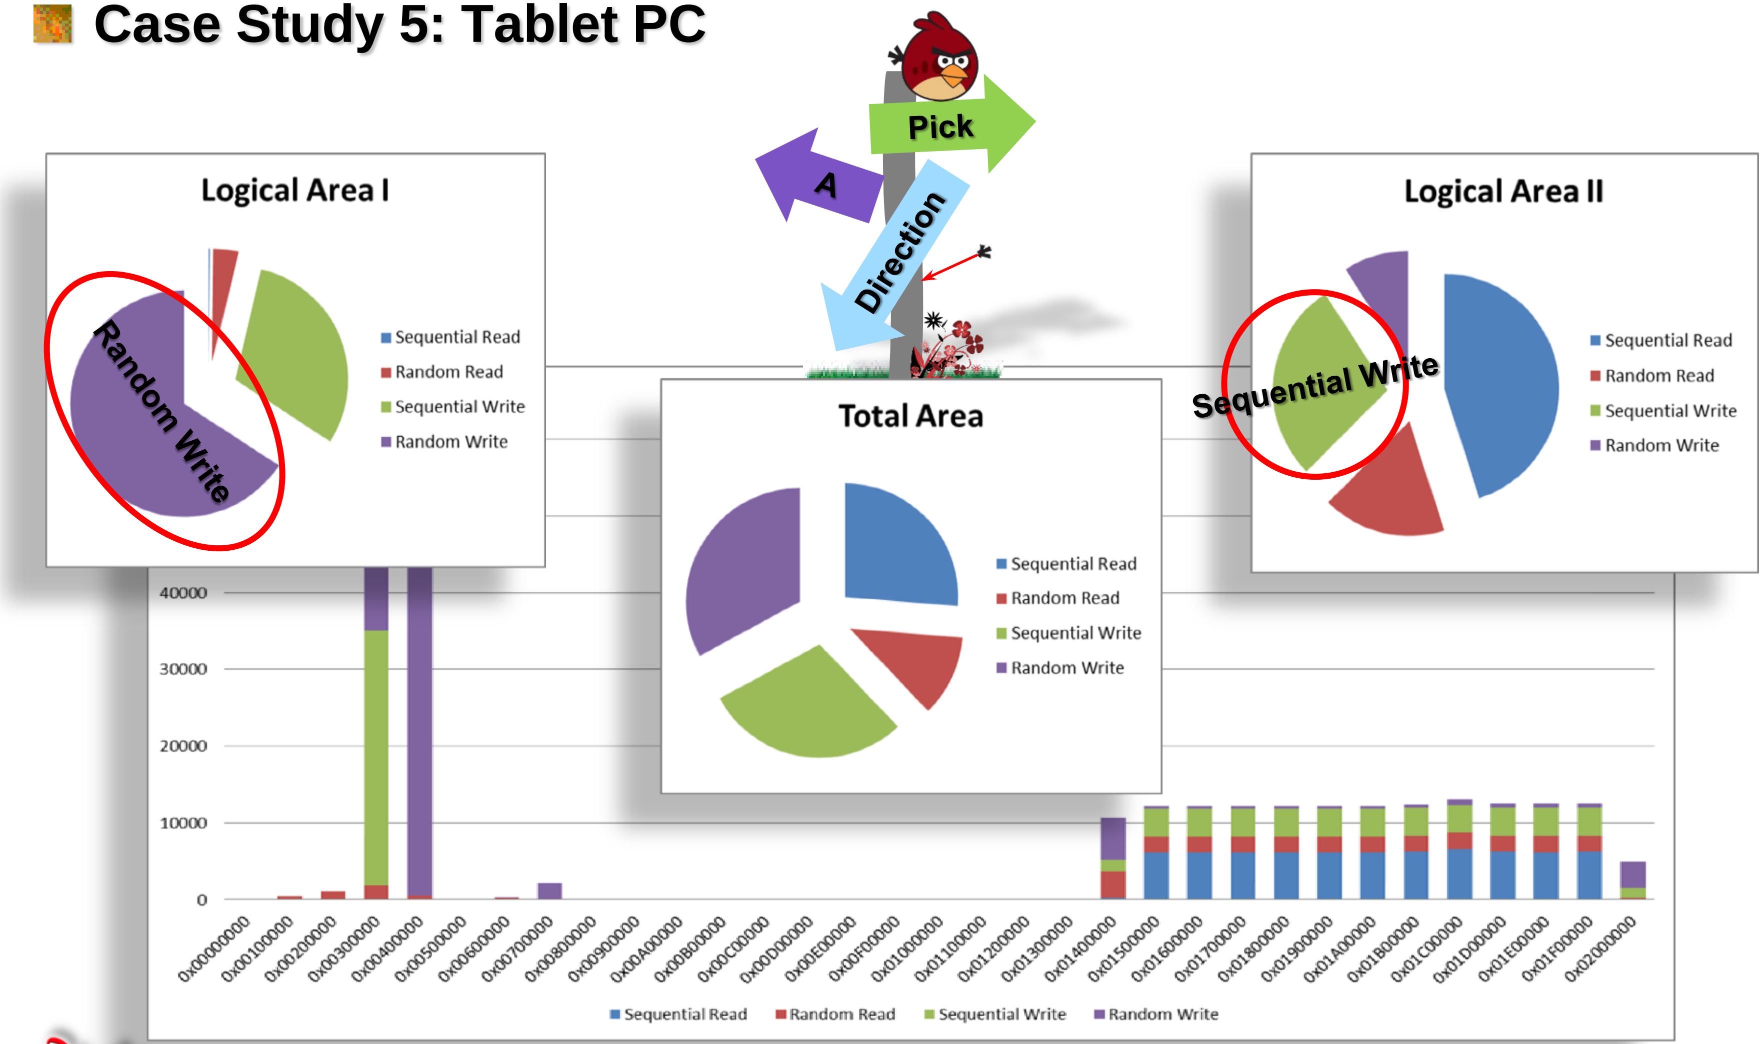

How does the future look like?

- **"Lynch" or "Lunch" ?**

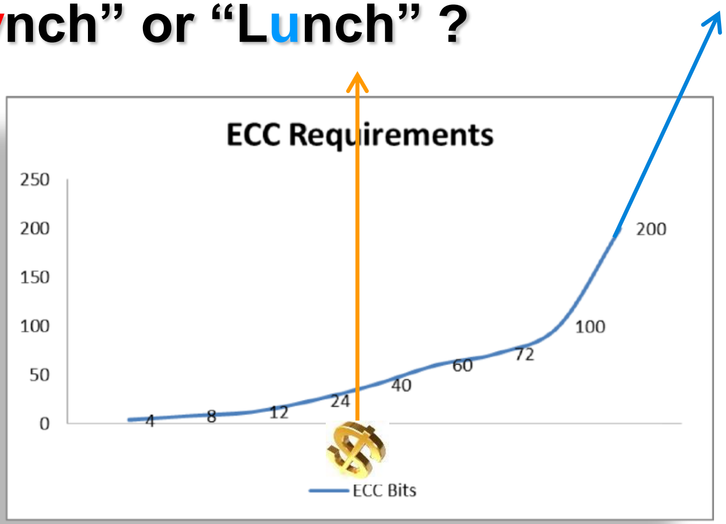

ECC Bits = F (WA, Target Life Span, Target UBER, Type of Material, …)

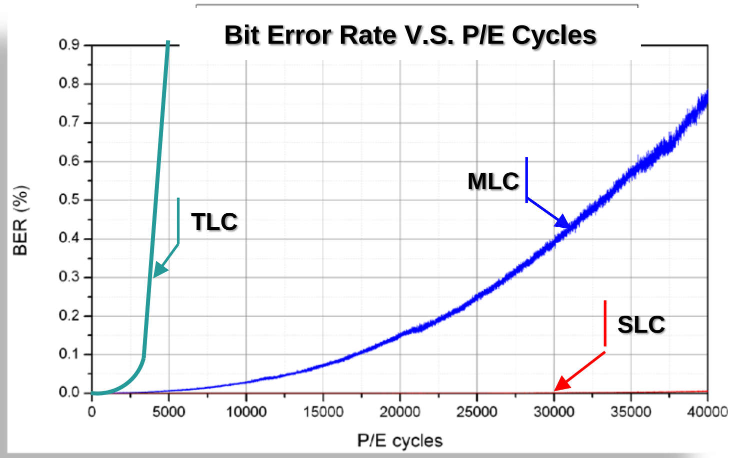

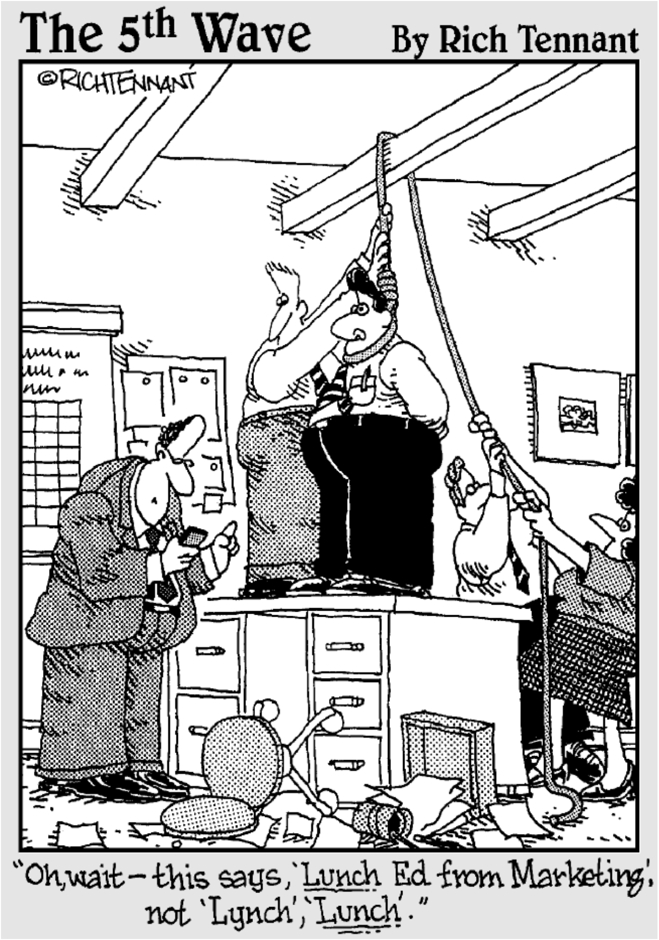

Source: "Managing Your Business With Outlook 2003 for Dummies" by Marcelo Thalenberg, Wiley Publishing, Inc.

How does the future look like?

- **Flash is more like a time-varying device ever.**

"Scott, when I said I wanted to grow old with you, I meant *gradually!*"

Source: http://www.glasbergen.com

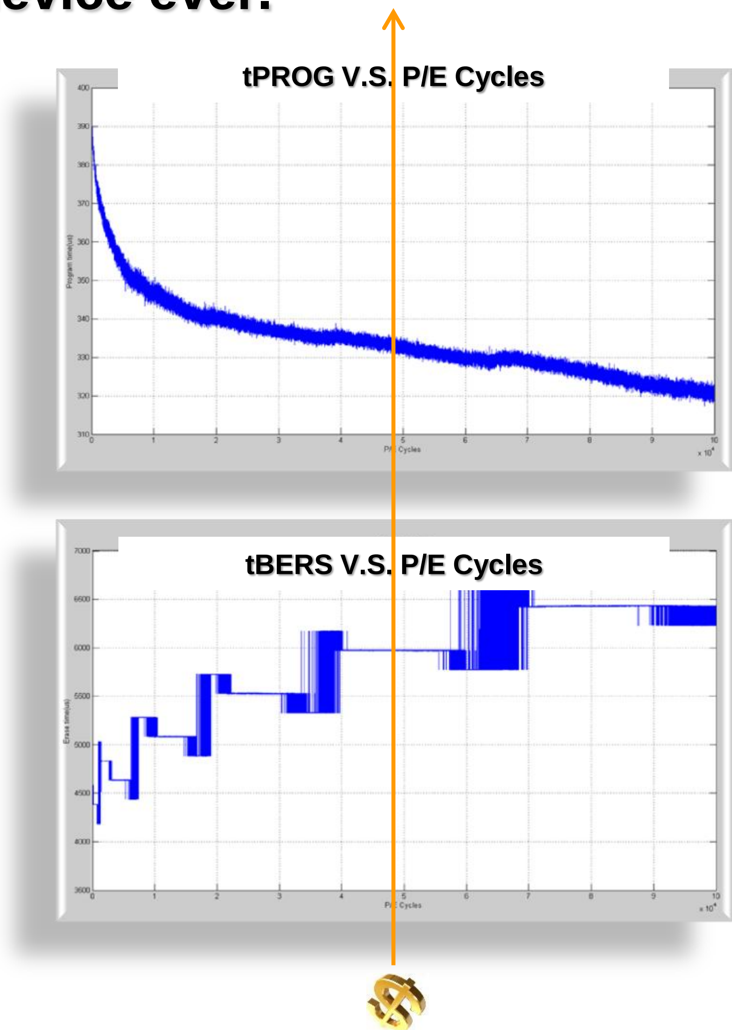

How does the future look like?

- **Should SSD's work harder to fit file systems or file systems be changed to be more flash-aware?** **(Maybe SSD guys do not need to work so hard.)**

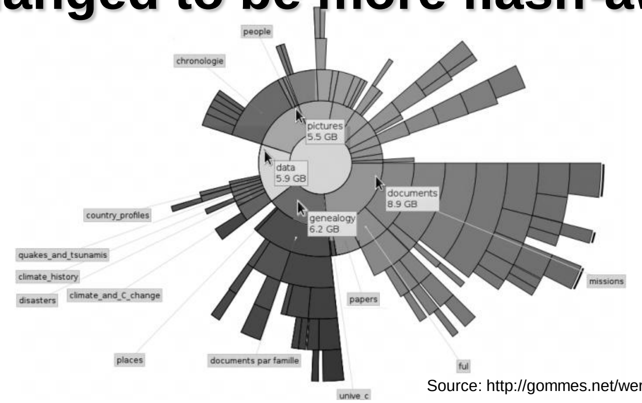

Source: http://gommes.net/wergosum/?p=674

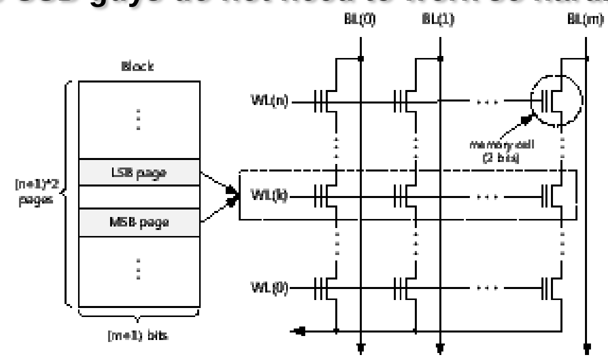

Source: Samsung FlexFS

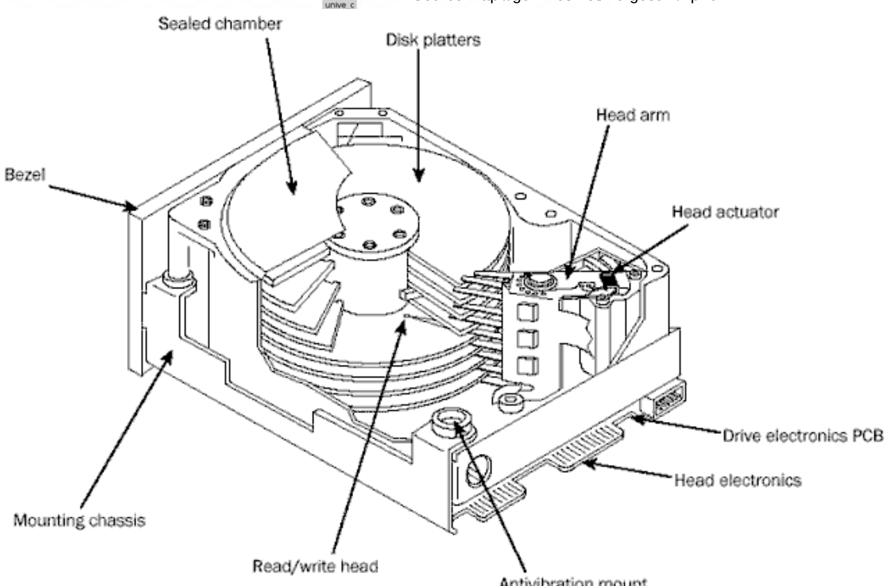

Source: http://alasir.com/books/hards/005-007.html

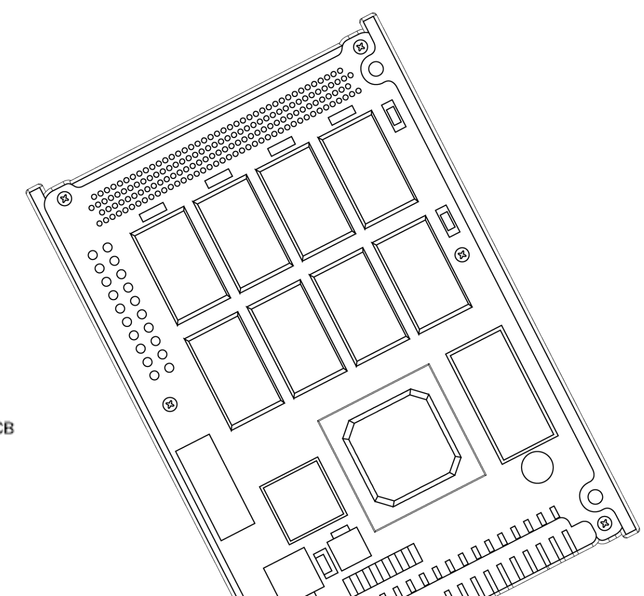

How does the future look like?

- **As a future SSD developer …** (As a buyer, don't over-stuff yourself.)

THANK YOU

and have a nice SSD