Optimize User Experience with Right Flash Storage Solutions

Young-joon Choi, CTO

AData Technology

2012 System Trend

- Predominantly Intel for Ultrabook and Tablet
- WindowsRT delay would push out eMMC
- SSD mSATA will keep dominant in client storage

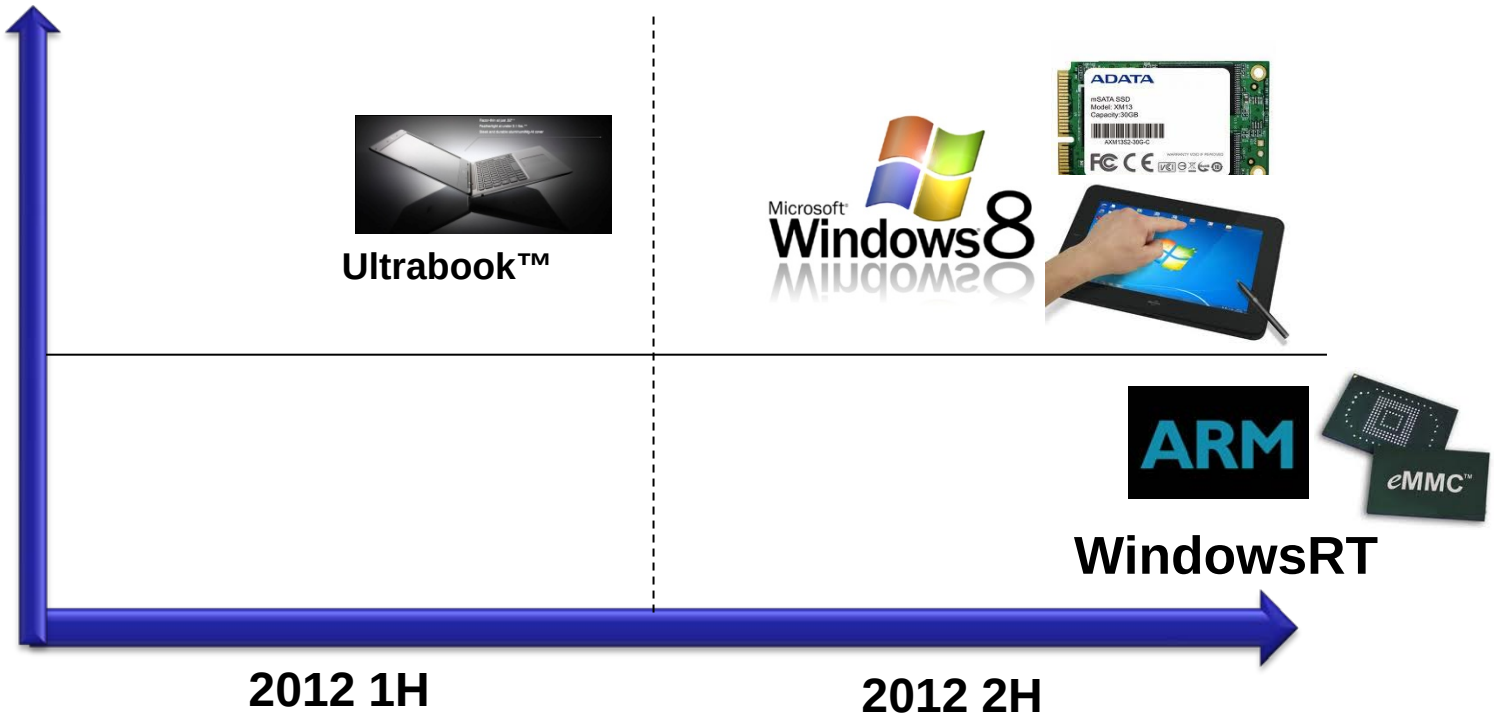

AFR of Intel SSD : ~0.6%

- Intel uses ComputeNAND with 0.1% die AFR
- But in SSD level, it is amplified to much higher

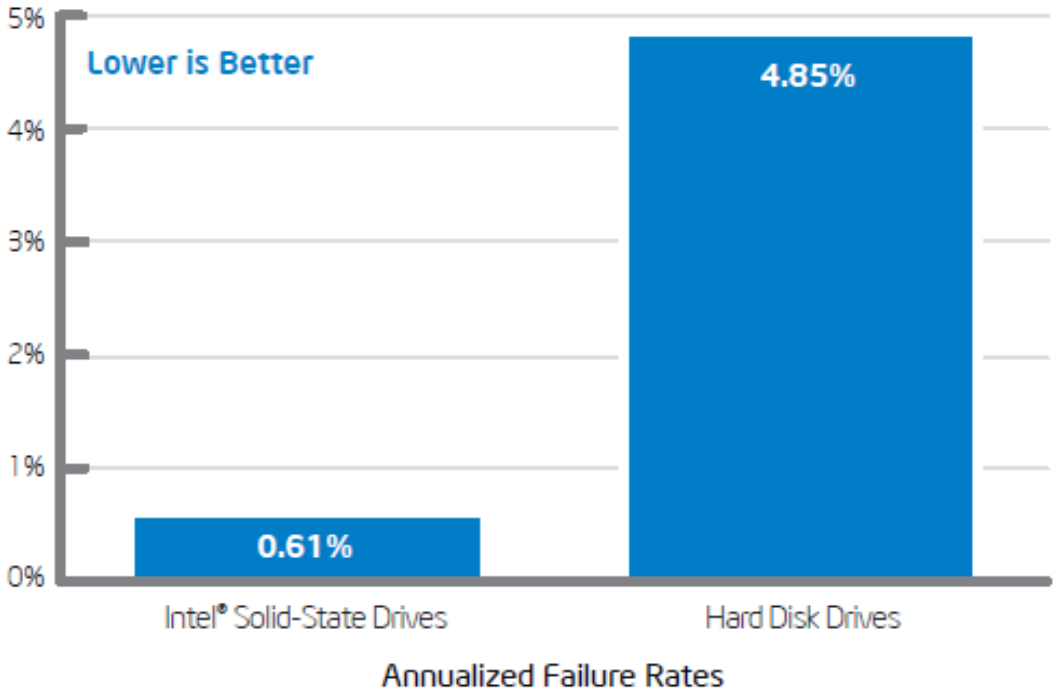

Figure 1. Analysis of Intel® Solid-State Drive failure data shows an 87 percent reduction in the annualized failure rate compared to hard disk drives, directly reducing IT total cost of ownership. Intel internal measurements, 2008 and February 2011.

Is SSD Reliable? Simple Math Says "Unlikely"

- Failure rates are multiplied proportionately to capacity
- When combined with other problems, it can easily exceed 1%

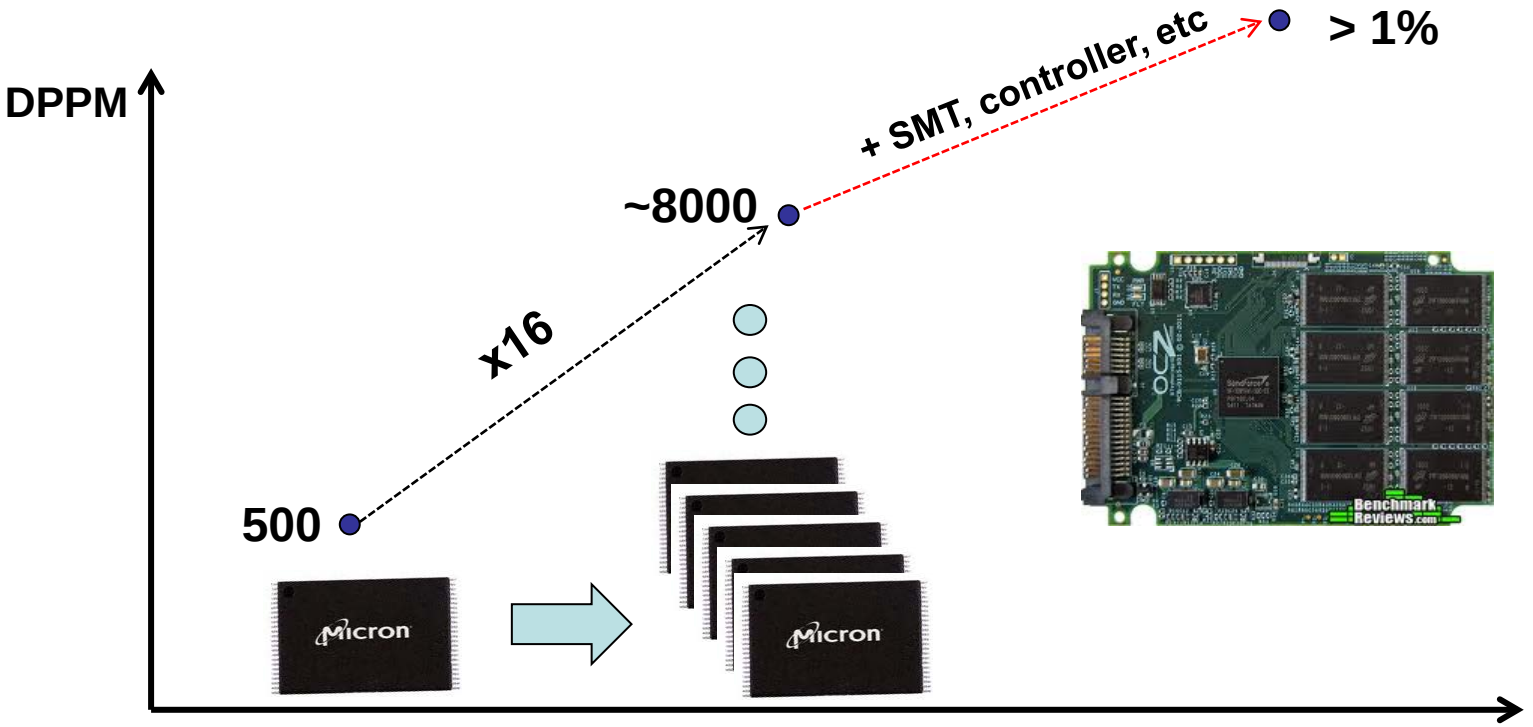

All SSDs Are Created Equal ?

- Eight SSDs with the same controller and the spec

Technical Details

. SandForce SF-2200 series SSD processor paired with qualified MLC NAND flash for best performance, value and reliability
. Max Read: up to 525MB/s, Max Write: up to 280MB/s, 4KB Random Write: 50,000(IOPS)
. SATA 6Gb/s,
. DuraClass technology
. DuraWrite extends the endurance

ADATA

All SSDs Are Not Created Equal

- Failure rate of some SSD is said to be as high as that of HDD

- **Rating: Garbage**

  Cons: drive failed within a matter of 12 days from installation - I am absolutely disgusted

  Pros: Inexpensive

- **Rating: unusable**

  Cons: ... Unfortunately, the old adage "you get what you pay for" rings true. The drive caused the computer to lock up while installing Windows the first time. Tonight (after for about 3 weeks) it caused BSOD several times and now I am stuck at the "Starting Windows" screen.

  Other Thoughts: Save your money and buy an better brand.

- **Rating : Great Price, bought 2, dog slow on z77**

  Cons: can't get past 12.5k/7.25k on any sata2/3 port in ahci mode, can't break 25k/10k when raid0 is configured.

- **Rating : absolute GARBAGE - failed in 18 days**

  Cons: failed within 18 days of installation along with my brand new OS, apps and files - what a headache!!! first time in 22 years that I had any drive fail in any system - I am so disgusted ...

  Other Thoughts: if you currently have this SSD - BACKUP YOUR FILES IMMEDIATELY!!! if you don't have it don't even think about getting it

Consumer-grade NAND Quality

Flash Qty Sorted	Bad Blocks Identified	Estimated DPPM of 128GB SSD
2420	3 blocks in 2 Flash (uECC x 2, High BER x 1)	~12,000 DPPM

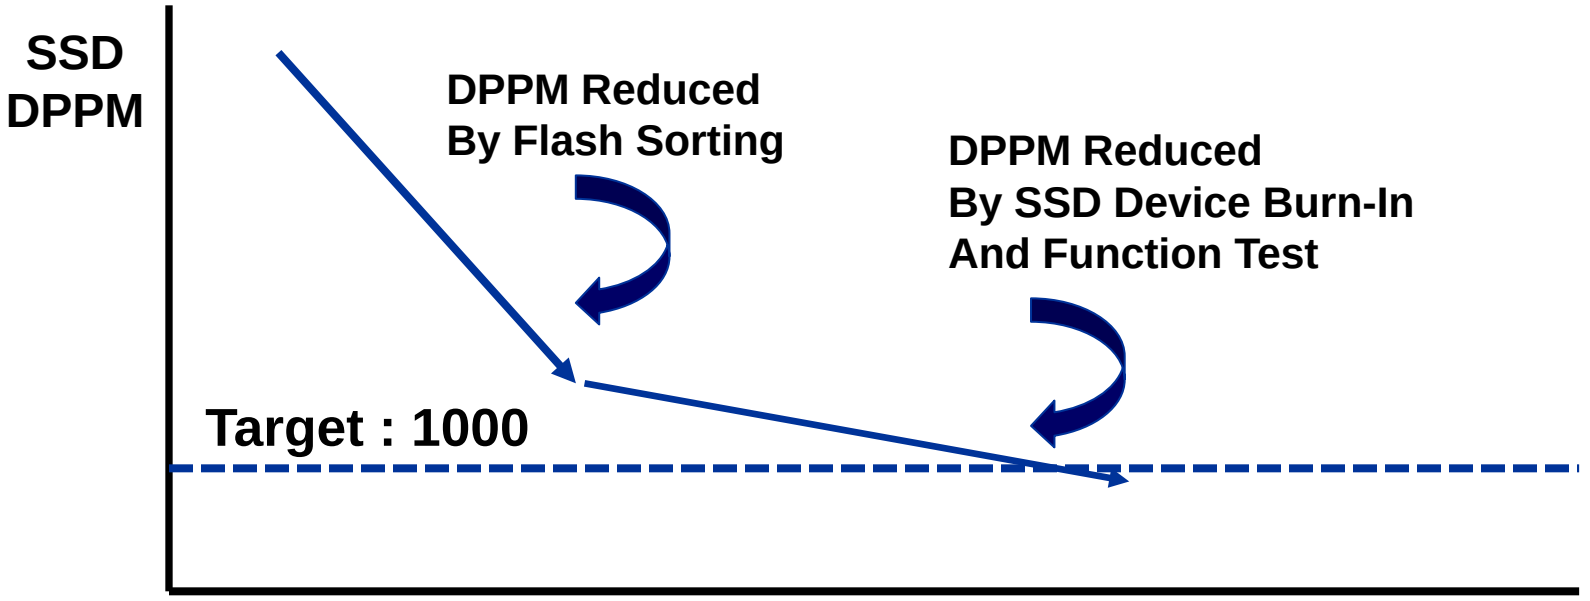

Server SSD : Challenges and Chances

- No power-failure protection needed → SATA-SSD
- RAID-controller compatibility → LSI RAID and controller
- Failure rate multiplied by no. of RAID → Very Low DPPM

Direct-Attached-Storage

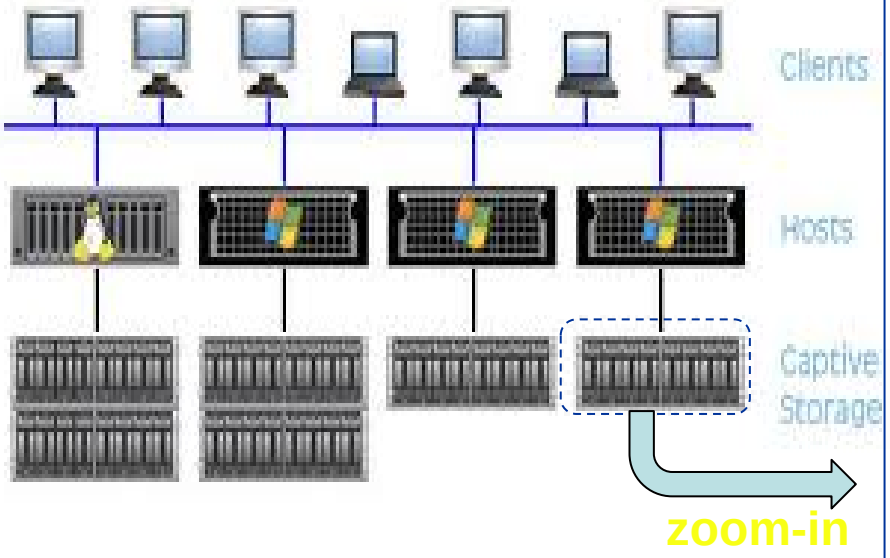

up to 10xSSD RAID

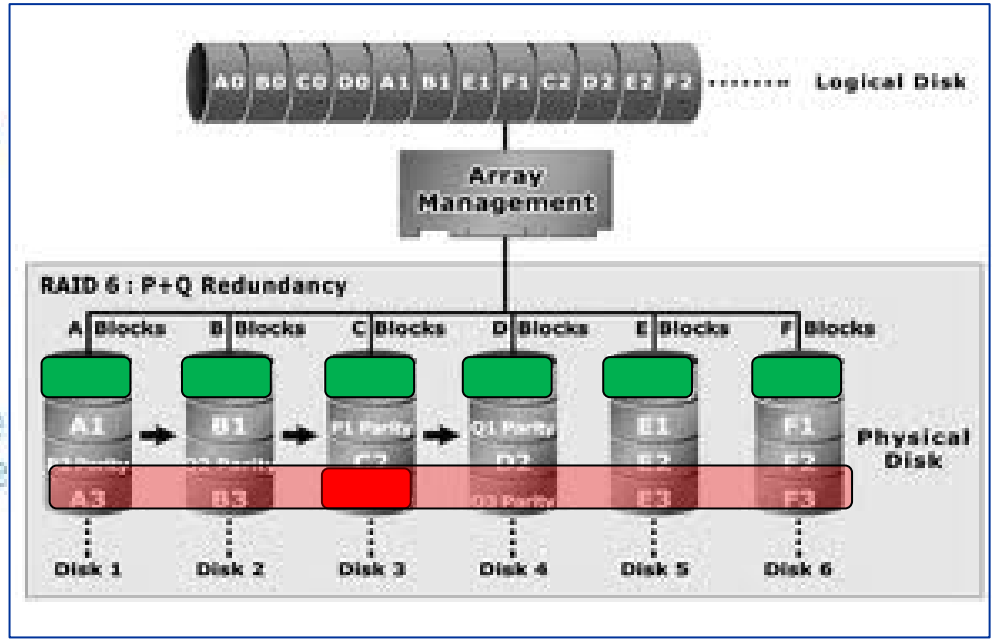

<1000 DPPM Quality and <0.5% AFR

- SSD production is done at Adata factory in Suzho,China
- Test history is written to every Flash component for traceability
- New products are validated through high volume quality validation process

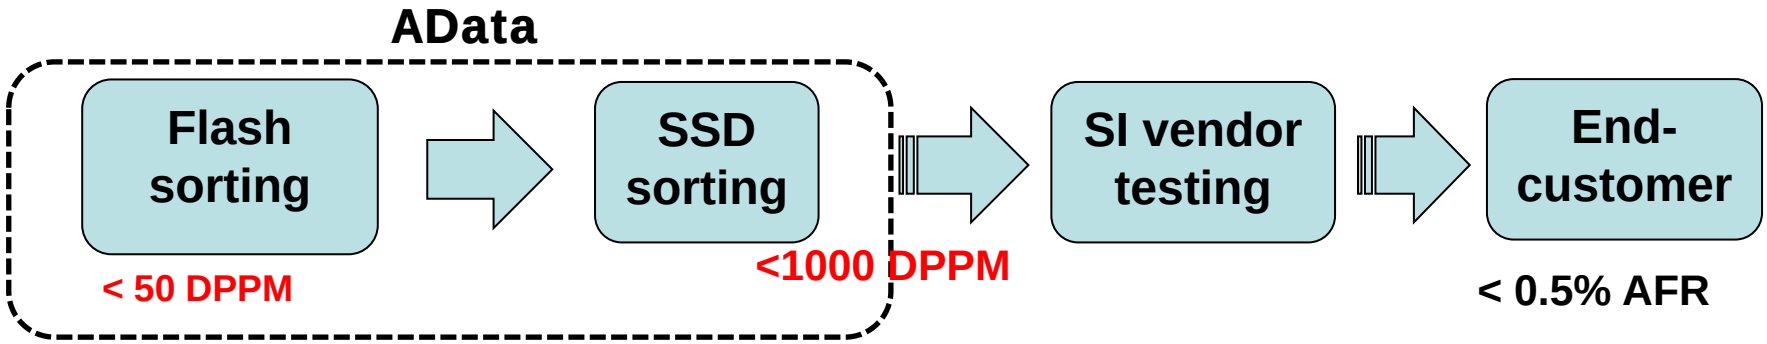

Internal RAID Suppresses Flash-related Failure

- Page-level ECC cannot protect gross W/L failures from W/L defect and W/L-shorts
- RAID5-type gives a good protection with 7% of capacity-expense

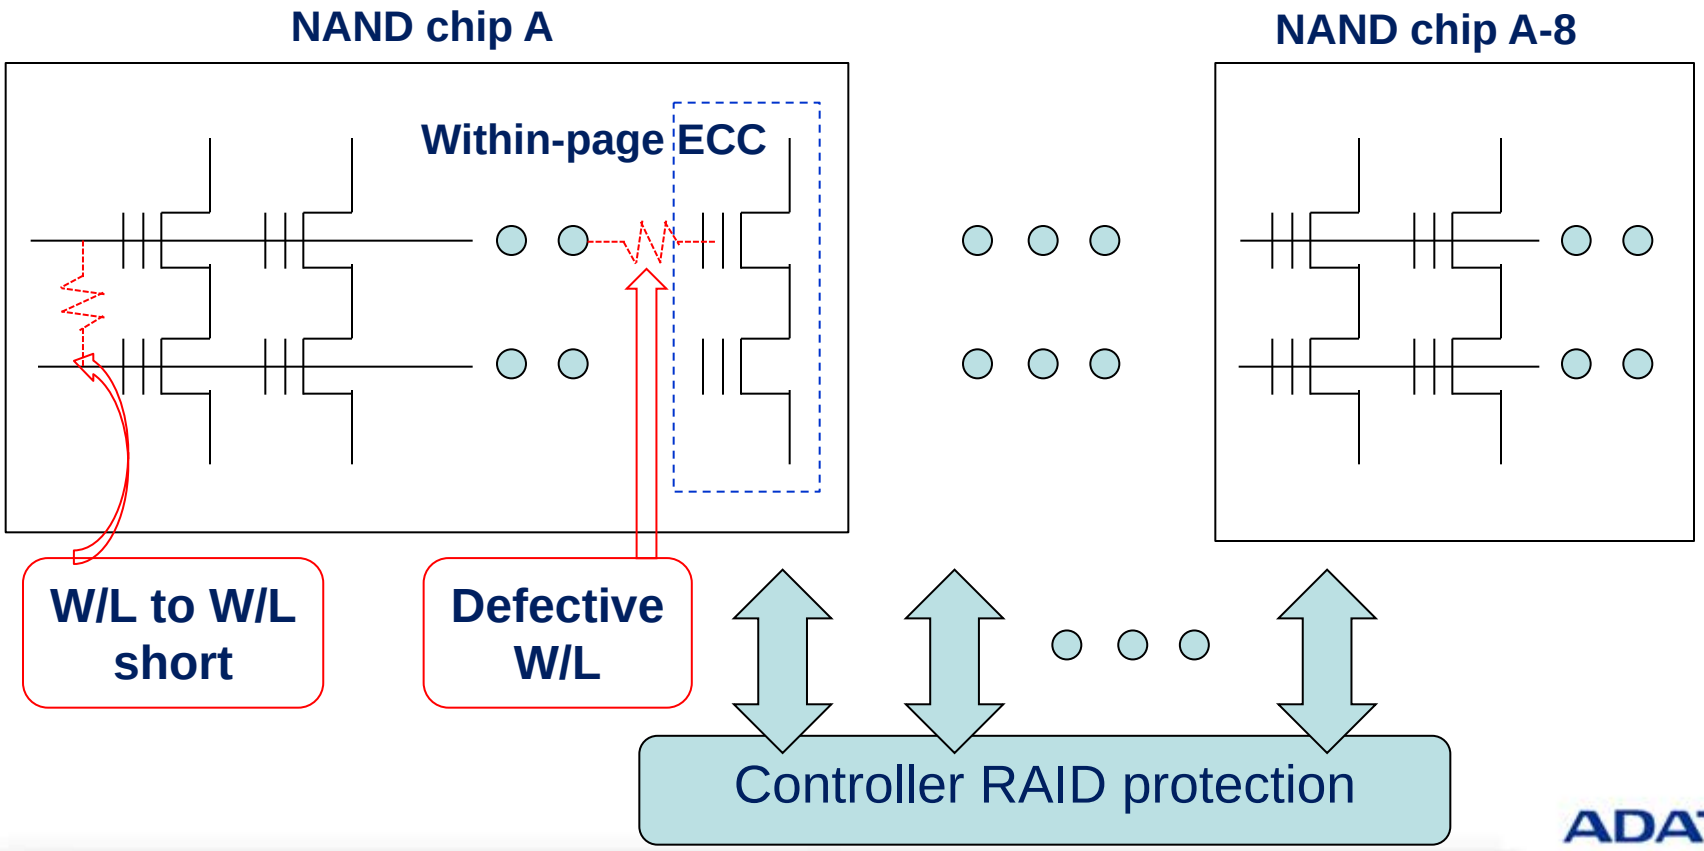

Performance : Compression?

- Compression saturates SATA bus and incompressible data transfer limited by Flash

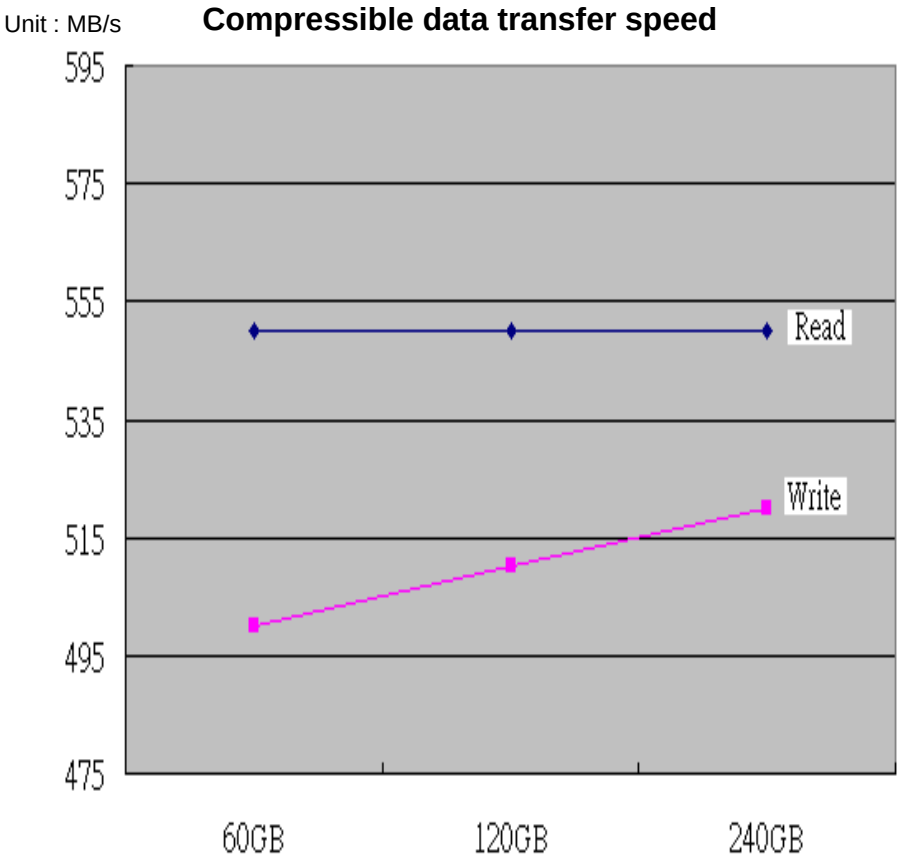

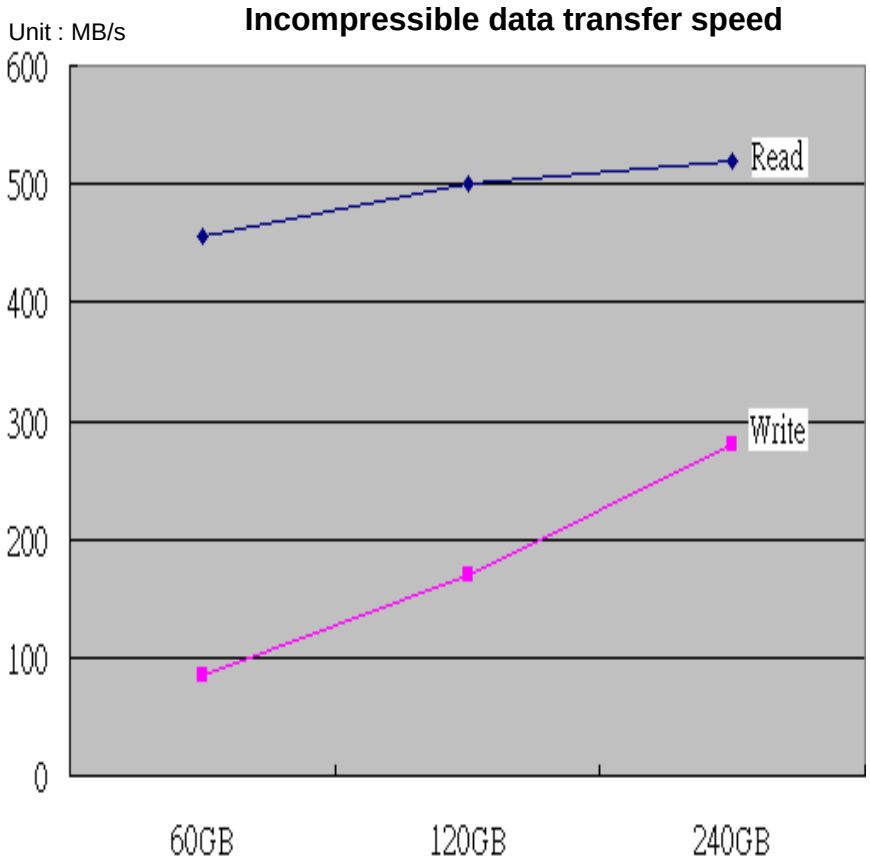

Compressible data transfer measure by ATTO and incompressible data transfer by AS-SSD

CDM & PCMark : Adata vs. C400

- PCMark score is consistent and outdoes C400
- C400 excels in sequential write and Adata excels in random write

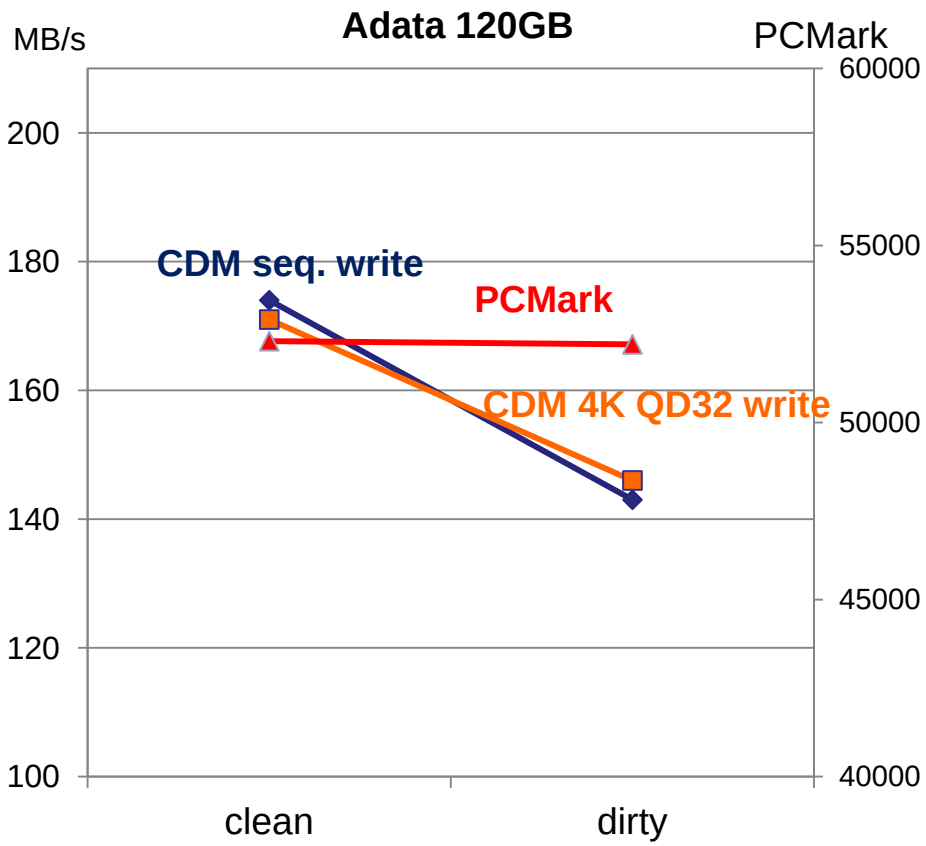

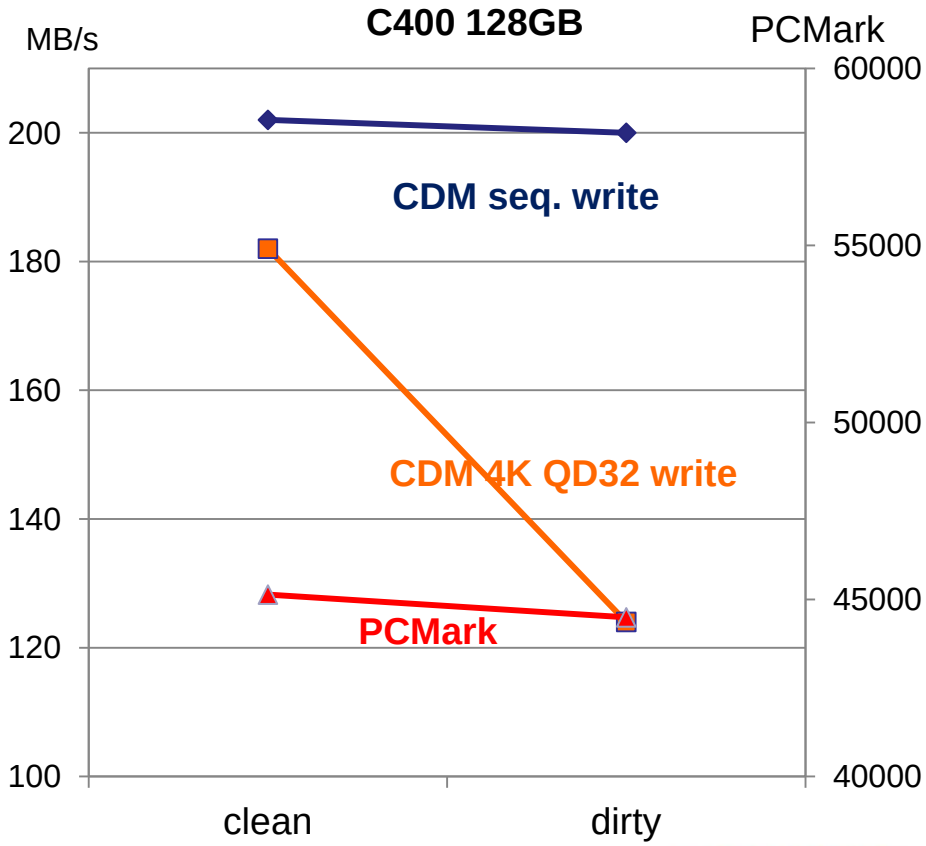

Chipset : HM65

Architecture Differences

- Adata : Page mapping scheme
  - Provides good random write performance
  - Compression provides less frequent garbage collection
  - Employs a foreground garbage collection (Just-in-time as needed)
    - Extends device life time by reducing early copy and erase
    - During GC, performance is dropped more than hybrid-mapping
    - As GC progresses, performance is recovered
  - Sequential performance suffers when there is no free space for GC
    - In real applications, trim prevents it from happening

- C400 : Hybrid-mapping scheme
  - Provides a good sequential write performance
  - Random write performance suffers more at dirty stage

Performance Benchmarking in Multitasking

- Performance of Adata 128GB is 20% better than Micron C400

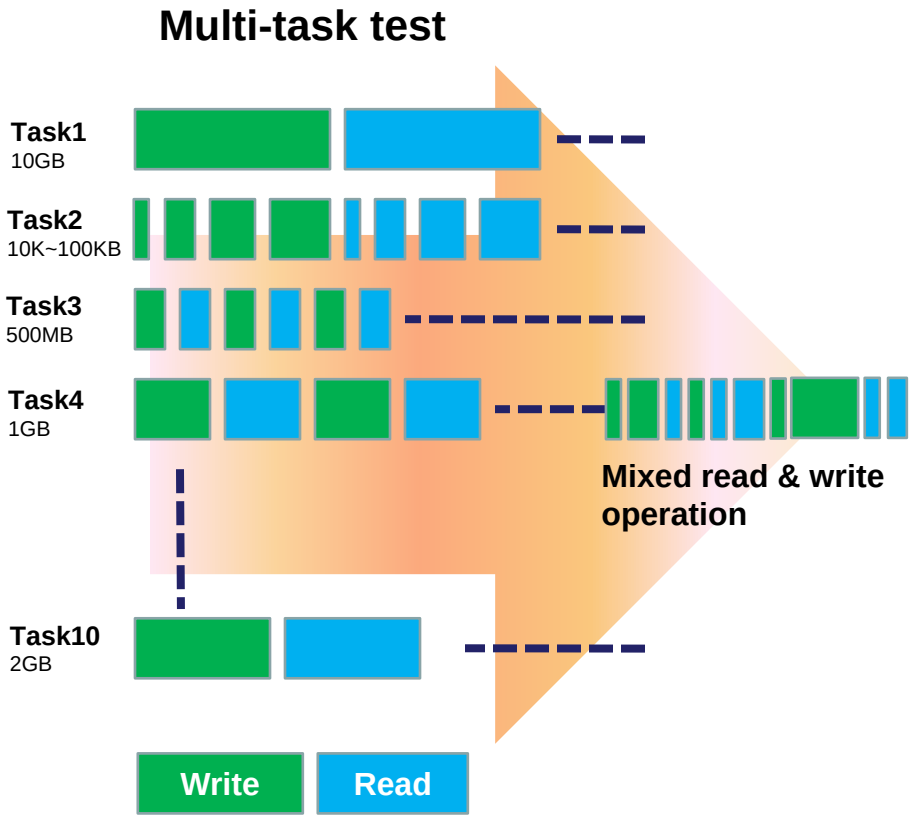

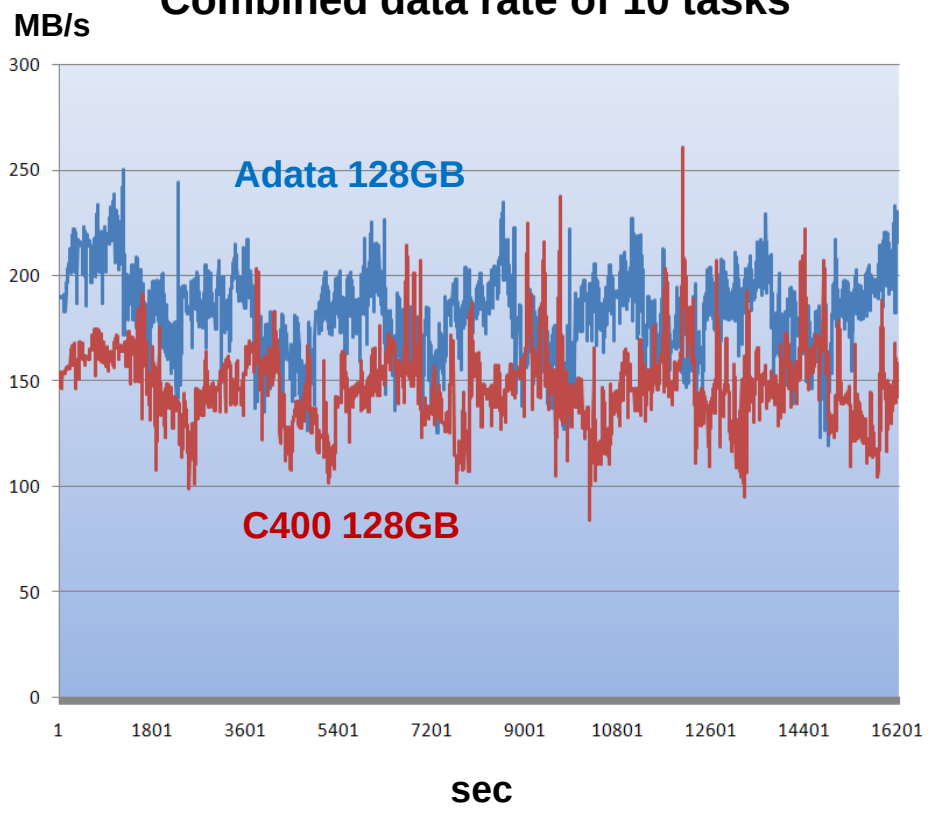

Conclusions of Performance

- Page-mapping and hybrid-mapping have both pros and cons
- Adata delivers higher peak and nominal performance than C400 which is faster in sequential writes
  - Real-time compression gives a further boost in performance and longer lifetime with lower failure rates
- Real-computing suggests sequential write tends to be random
  - Analysis showed Adata SSD with 20% faster than C400

Windows7 Boot Process Analysis (I)

- Read at SMSSInit occupies the majority of data transfers and 2/3 of total I/O delays in Windows boot sequence

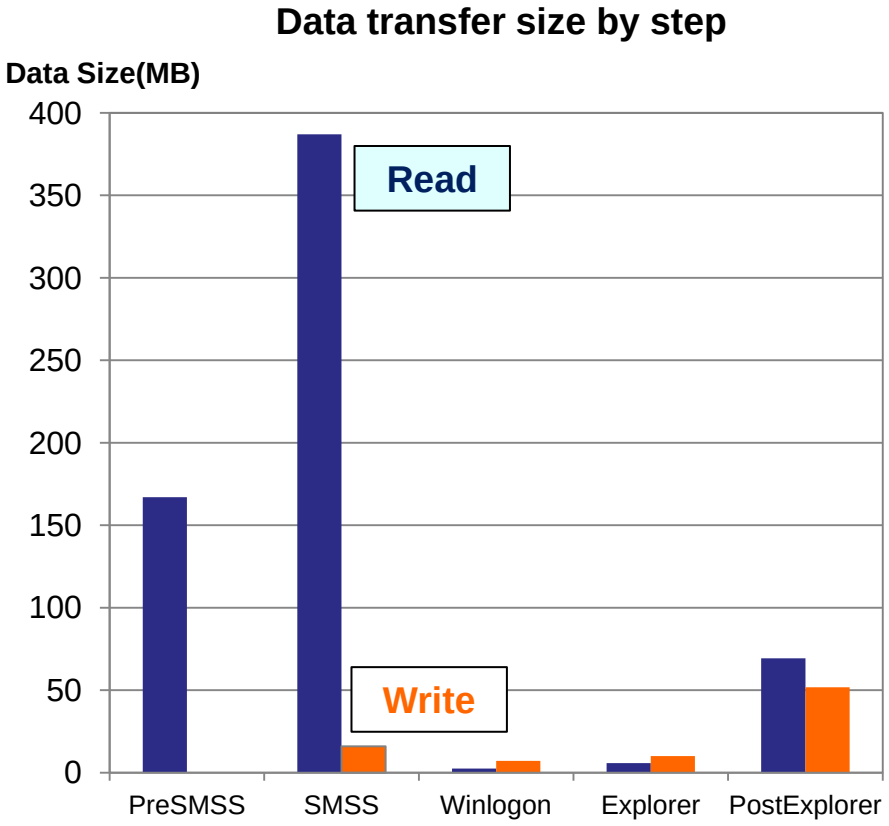

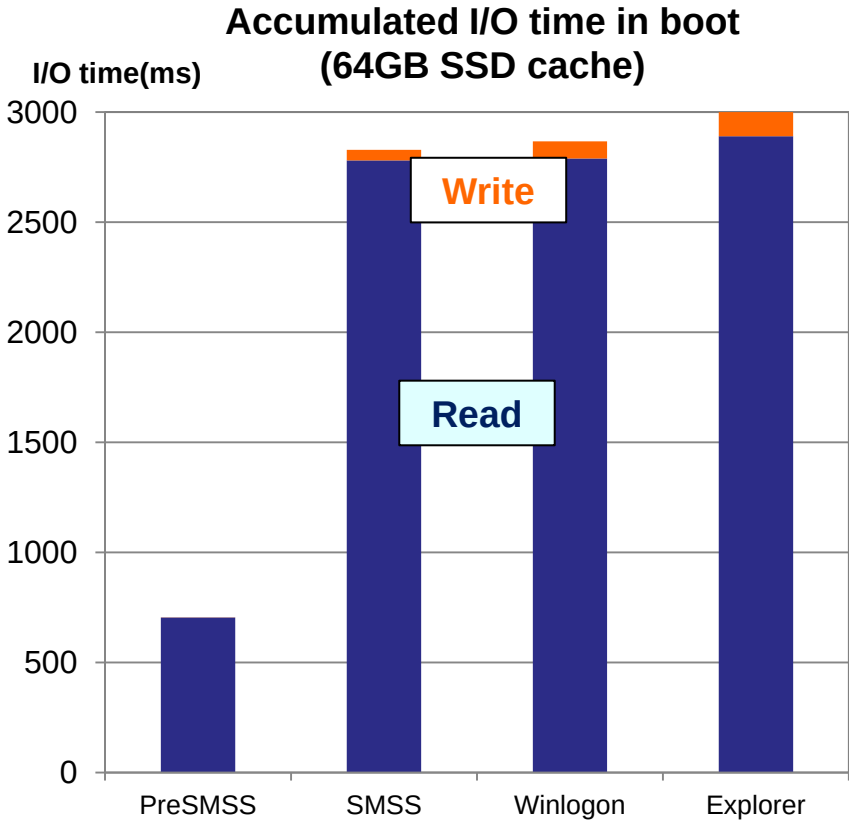

Analysis by Windows Performance Analyzer from Microsoft

ADATA

Windows7 Boot Process Analysis (II)

- 4-channel cache SSD gives similar transfer rate to 8-channel cache
- Compressed boot data boosts the performance over raw data rate

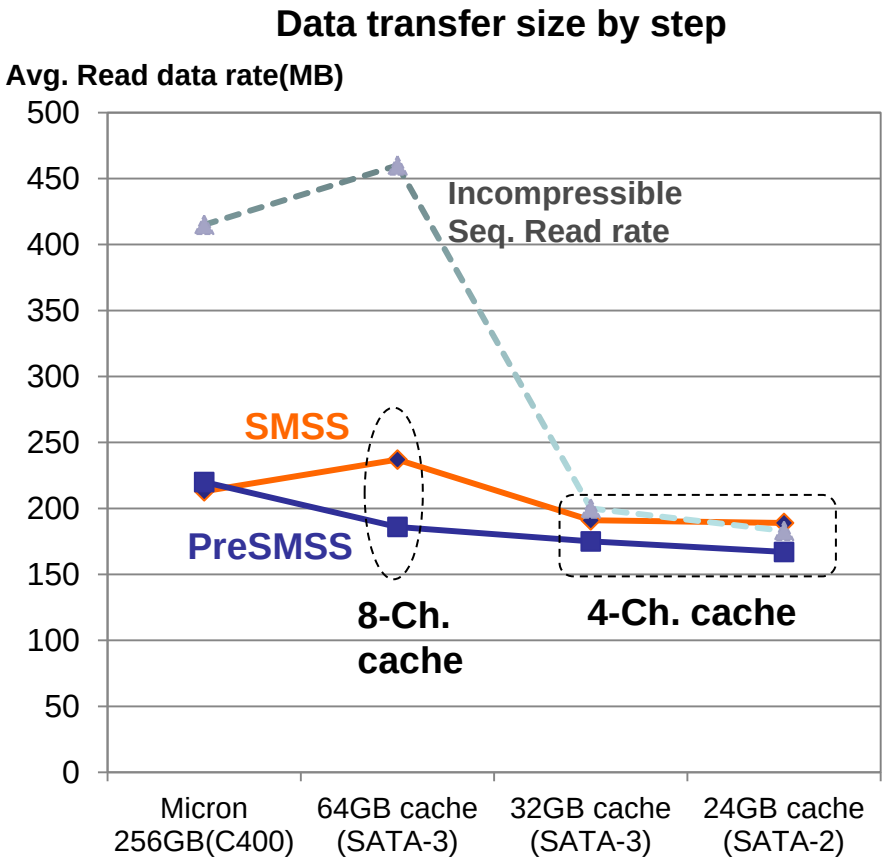

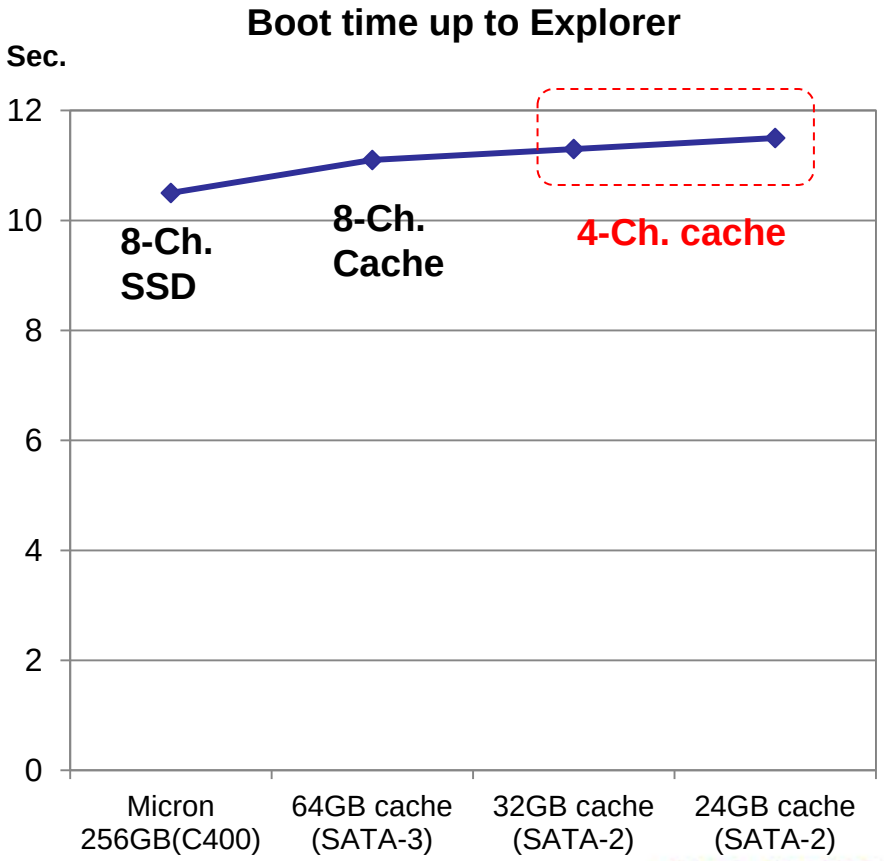

Key Take-Away

- I/O operations during Windows boot are read-centric
- SMSSInit occupies the major portion of boot sequence
- 3 and 4-channel cache SSD delivers comparable performance to 8-channel cache despite less than half of sequential read rate for incompressible data
- Compression should have pushed read rate of cache-SSD over 2X at boot since boot image is mostly compressible
- The high read rate of SSD cache was achieved at three times of booting by saturated high hit ratio
- Compression is also effective in most of PC applications, especially for virtual memory page-file swapping, office applications and internet browsing

Cache-SSD Wear Varies by File Types

- For large sequential data, write bypasses SSD-cache into HDD
- Random and small data are written to SSD-cache
- Multi-media files tend to be bypassed into HDD

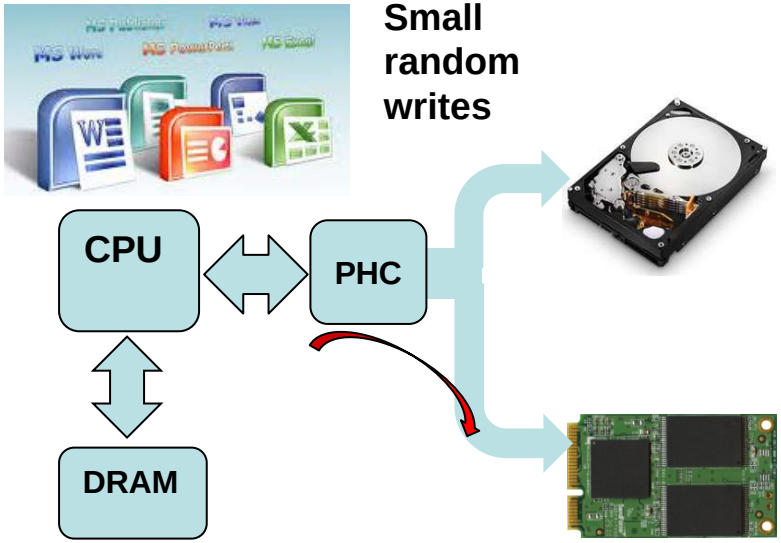

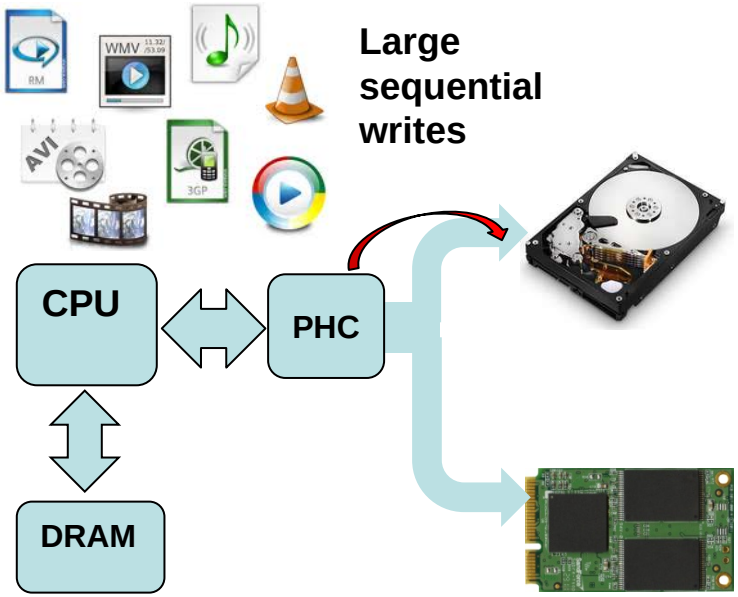

Life Time Estimation of Cache-SSD

WAF	O/S Layer	SSD Cache without Compressor	SSD Cache with Compressor
Large files	Low	Low	Low
Small files	High	High	Low

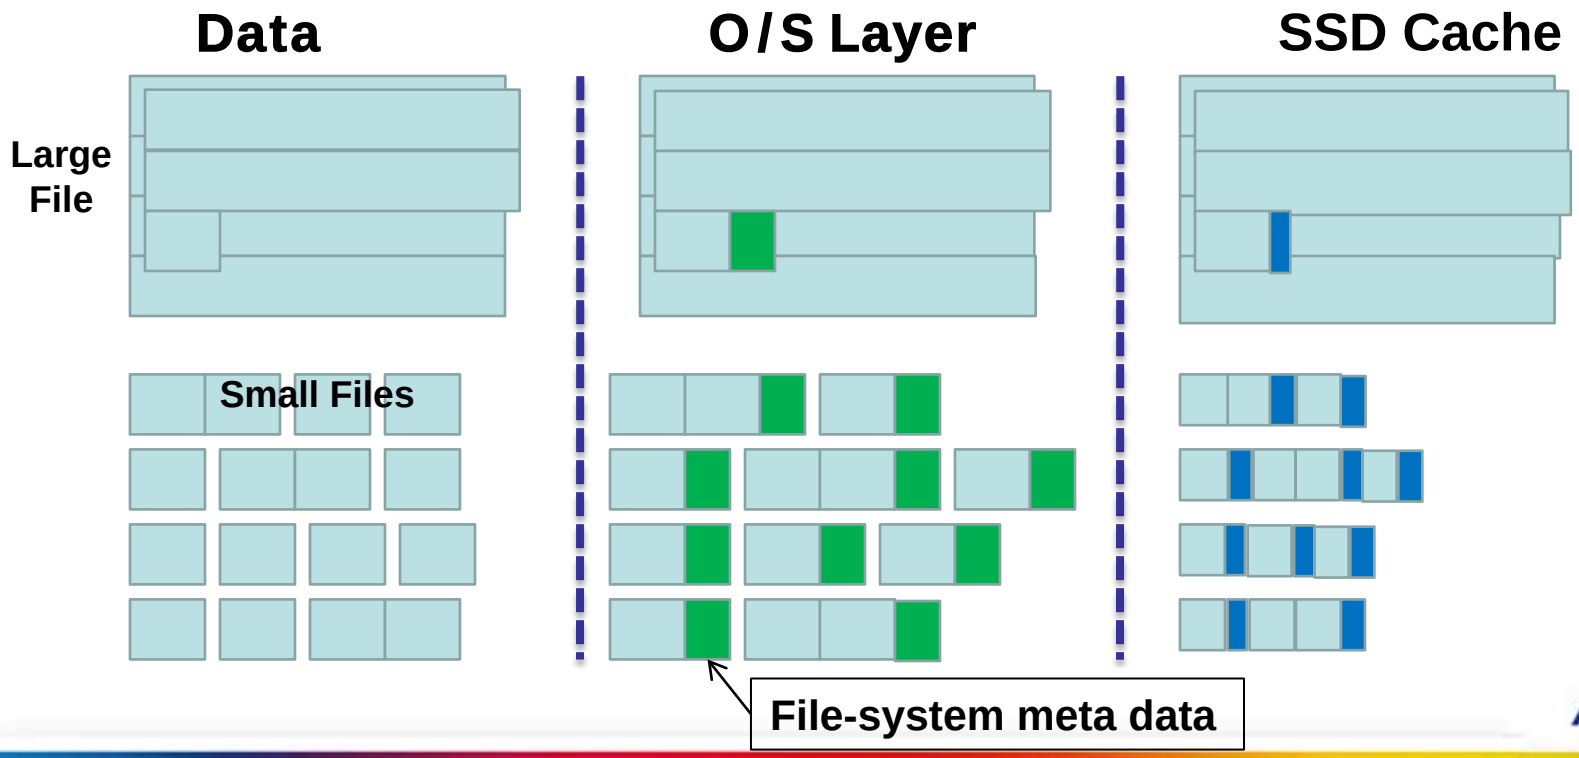

Estimation of Life Time

- Compression balances endurance of small random writes

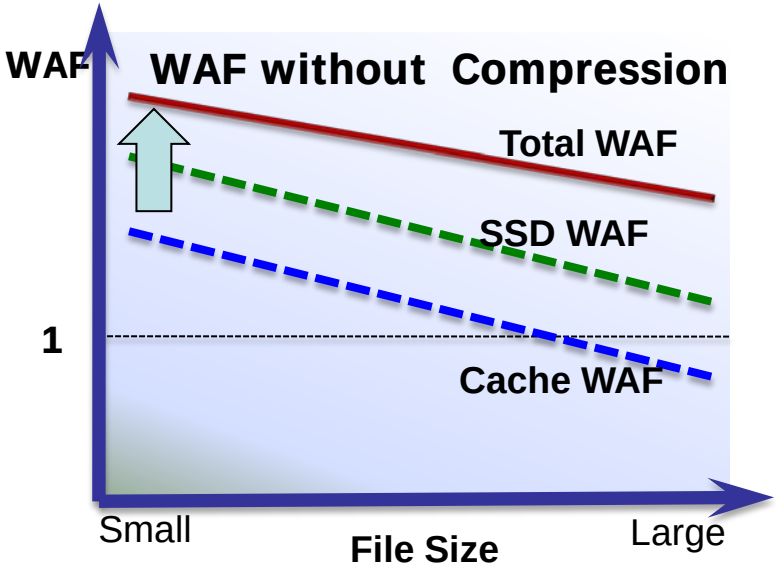

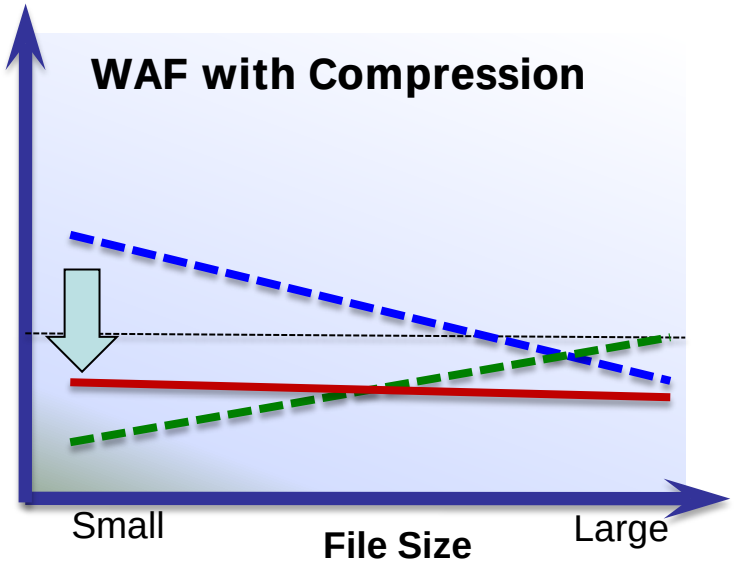

iRST	Small Files (1GB)	500MB	1GB	2GB
Total Write Capacity(GB)	**55**	**25**	**50**	**100**
Total Flash Write	41	18	30	65
Total SATA Write	77	18	28	78
Cache WAF (SATA Write / Total write)	**1.4**	**0.72**	**0.56**	**0.78**
SSD - WAF	**0.53**	**1.00**	**1.07**	**0.83**
Total - WAF	**0.75**	**0.72**	**0.60**	**0.65**

24GB m-SATA Feasibility

- Meets Ultrabook performance criteria with reasonable reliability
- Sub-1 WAF from on-chip compression extends reliability greatly

Item	Item name	Enable Intel Smart Response/ Maximized mode(full size)	
		HDD + XM13 22GB (3CH)	HDD + XM13 30GB (4CH)
1	S4-Hibernate (Intel Spec : Resume 7 sec)	1.2 Sec.	1.2 Sec.
2	PCMark Score (Intel Spec : 16,000)	22,102	22,765
3	Video Editing using Windows Movie Maker (Intel Spec : 80 MB/s)	90 MB/s	92 MB/s

Scenario	Measured host write	WAF*note1	Flash write
Daily workload	~10GB (MM'07 @8hr)	0.8	8GB
Hibernation	~6GB (10 times)	1	6GB

Daily total : 14GB

- **Life time of 24GB : 24GB * 3K PE cycles / 14 =~ 14 Years**

*** Note 1 : Write Amplification Factor was taken from Smart parameter**

High Capacity mSATA Roadmap

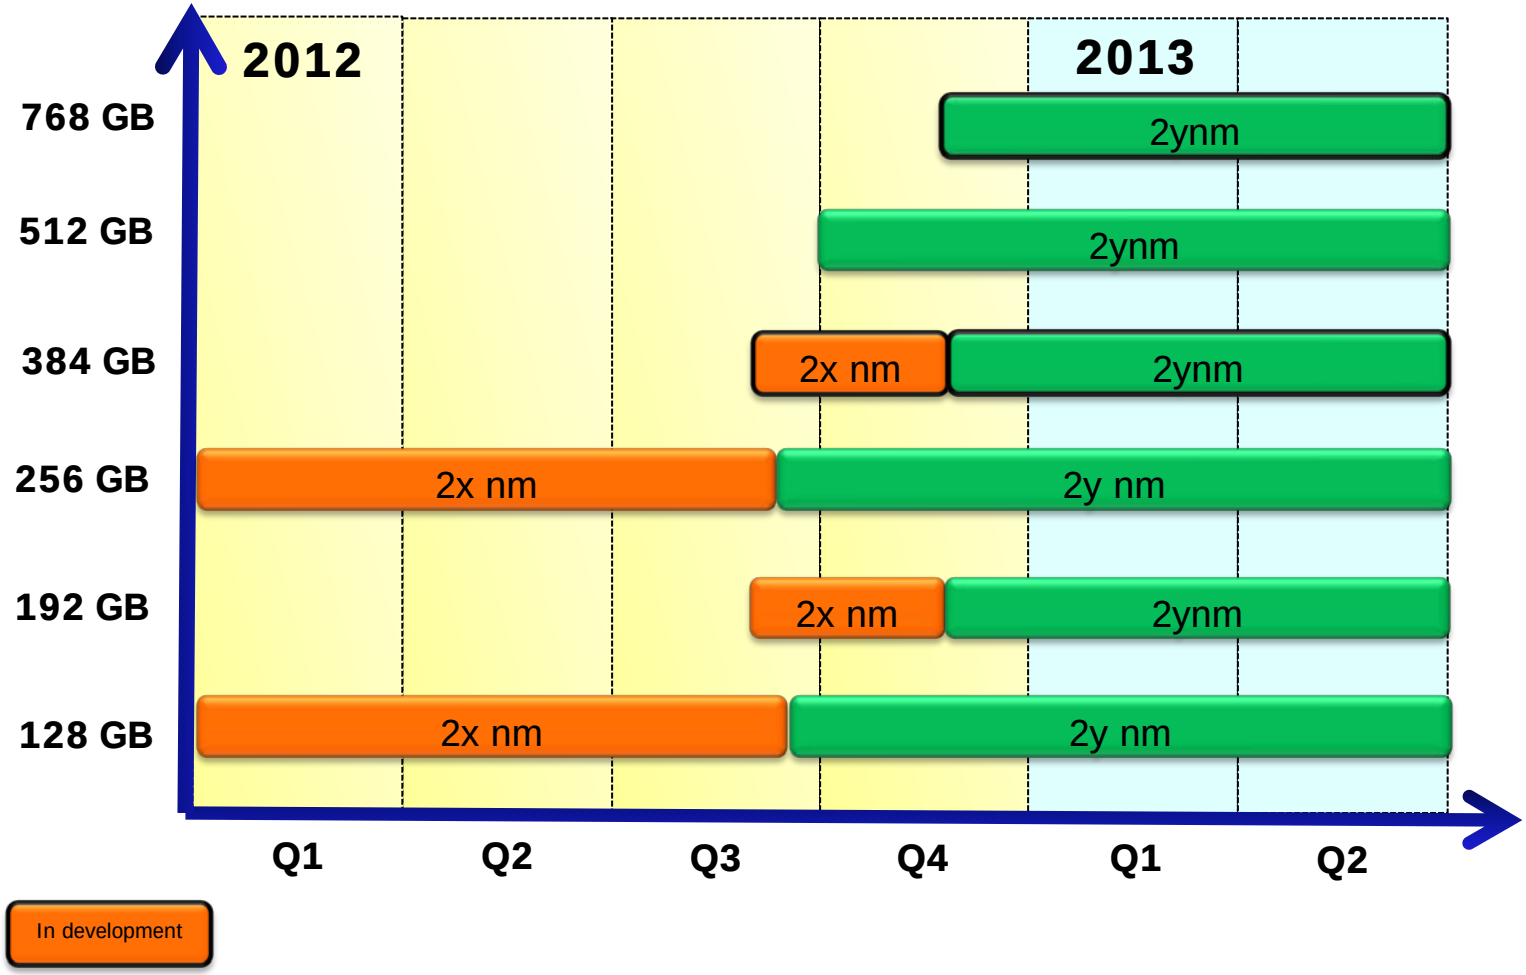

- Support 384/512/768GB of mSATA solution in 2012.Q4 to fulfill the high capacity request of SSD.

Conclusions

- In-house Flash testing capability is the key to achieving high quality SSD for <0.5% AFR
- In real-computing, sequential write tends to be random and ADATA SSD provide excellent random performance to get better user experience.
- Compression is effective in most of PC applications, with page-file swapping, office applications and internet browsing. It also lowers write-amplification, extending reliability for cache-SSD
- Adata provides leading high quality SSD solutions with strong product portfolio and in-depth system knowledge

THANK YOU