



## The Agenda

- Review of SSD technology trends
- Windows 7 SSD enabling technology
  - Learning the new features of Windows 7 for SSD identify and Trim
  - The importance of NTFS partition alignment to SSD geometry, and the implementation
  - Proposed Windows 7 logo requirements related to SSD
- Remaining challenges
- Ongoing effort and progress

WinHEC  
2008

## SSD Performance Trends (Sequential write)

Sequential performance continues to improve

MLC drive's Performance is increasing

Sequential performance advantage is big and real



Source: a subset of sample data from internal lab

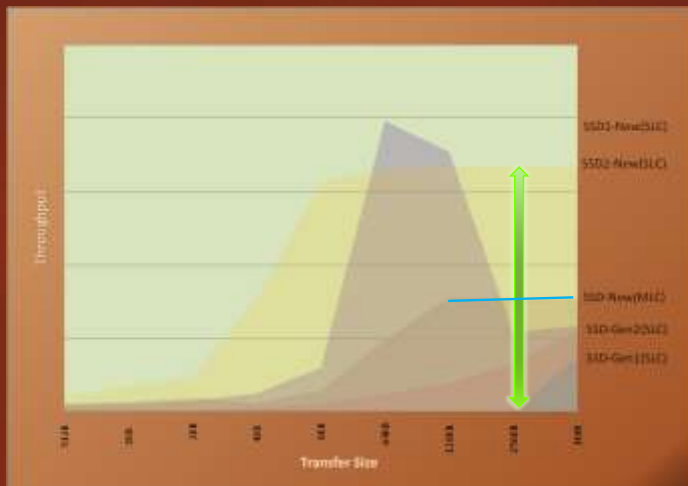
WinHEC  
2008

# SSD Performance Trends (Random write)

Random write speed is increasing

MLC drive's random write is also improving

Random performance issues are being solved



Source: a subset of sample data from internal lab



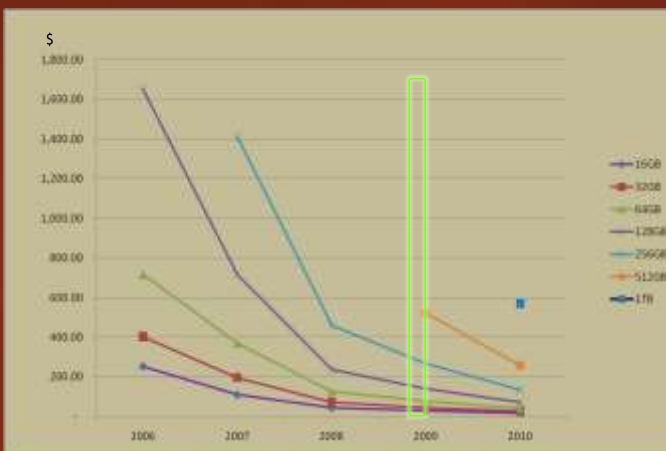
# SSD Cost Trends

The mark shows where the cost is today

And it continues down

The device cost will be in affordable range by 2010

1TB SSD is on the radar



Source: Semiconductor Forecast Worldwide--Forecast Database [SEQS-WW-DB-DATA], Gartner August 2008, by Joe Unsworth, et al.





# SSD Is On The Path To Its Full Potential

Microsoft WinHEC 2008



## Windows 7 SSD Enabling Technology

Microsoft WinHEC 2008

## Windows 7 Enables Improved Endurance for SSD Technology

- SSD can identify itself differently from HDD in ATA as defined by ATA8-ACS Identify Word 217: Nominal media rotation rate

7.16.7.77 Word 217: Nominal media rotation rate  
Word 217 indicates the nominal media rotation rate of the device and is defined in table 38.

Table 38 — Nominal Media Rotation Rate

Value	Description
0000h	Rate not reported
0001h	Non-rotating media (e.g., solid state device)
0002h-0400h	Reserved
0401h-FFFFh	Nominal media rotation rate in rotations per minute (rpm) (e.g., 7200 rpm = 1C20h)
FFFFh	Reserved

- Reporting non-rotating media will allow Windows 7 to set Defrag off as default; improving device endurance by reducing writes

WinHEC  
2008

## Windows 7 Enables Optimization for SSD Technology

- Microsoft implementation of “Trim” feature is supported in Windows 7
  - NTFS will send down delete notification to the device supporting “trim”
    - File system operations: Format, Delete, Truncate, Compression
    - OS internal processes: e.g., Snapshot, Volume Manager
- Three optimization opportunities for the device
  - Enhancing device wear leveling by eliminating merge operation for all deleted data blocks
  - Making early garbage collection possible for fast write
  - Keeping device’s unused storage area as much as possible; more room for device wear leveling.

WinHEC  
2008

## NTFS Partition Alignment is Important for SSD Performance

- The alignment of NTFS partition to SSD geometry is important for SSD performance in case of Windows XP and Windows XP upgrade to Windows Vista and Windows 7
  - The first Windows XP partition starts at sector #63; the middle of a SSD page
  - Misaligned partition can degrade device's performance down to 50% caused by read-modify-write
  - The example with 4k page size
  - Implementing correct alignment according to the latest ATA and SCSI spec.

WinHEC  
2008



## Proposed Windows 7 Logo Requirements For SSD

WinHEC  
2008

## Windows 7 Storage Logo Proposal (1-3)

- Proposed Windows 7 logo requirements related to SSD
  - Recommendation: If a device requires Windows defragmentation to be turned off by default, the device should report its Nominal Media Rotation Rate as 0001h Non-rotating media (e.g. solid state device) as per the ATA8-ACS1 specification, section 7.16.7.77.
  - Requirement: For all data transfer sizes up to 128KB, command completion time of a single read or write operation shall not exceed 5 seconds.

WinHEC  
2008

## Windows 7 Storage Logo Proposal (2-3)

- Requirement: For all data transfer sizes up to 128KB, and over any period of sustained I/Os of or exceeding 10 seconds, 98% (7 sigma) of Read I/O command completion time shall not exceed 3 seconds, and, the Read I/O command completion time coefficient of variation  $C_v = \frac{\sigma}{\mu}$  shall not exceed 0.5.
- Requirement: If the device implements “Trim” support, the “Trim” implementation shall comply with ATA8-ACS2 proposal e07154r6 (Data Set Management Commands Proposal for ATA8-ACS2) section 5.3 and section 6.2. The command completion time of Trim shall be less or equal to 20ms for every GB of the data range being trimmed, and Trim shall not add or remove LBAs from the NV Cache Pinned Set.

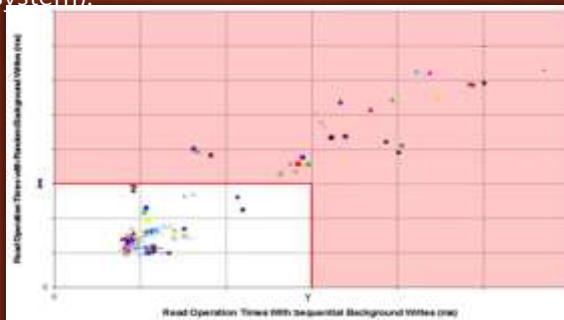
WinHEC  
2008

## Windows 7 Storage Logo Proposal (3-3)

Requirement: The Windows 7 Beta Windows System Assessment Tool (WinSAT) write policy flush test for the device must pass for any visible storage space utilization up to 95% (% of utilization as % of "used space" seen through the Windows file system).

Giving read a priority can be important when there is a long queue of writes

The result is better user experience of system responsiveness



WinHEC  
2008

## Challenges Ahead

- SSD performance and quality are scattered widely with product on the market
  - Maintaining a consistent good user experience
- SSD endurance is equal to the safety of user's data
- Data retention becomes more important
- The answer to questions
  - Does SSD has its own defrag issue?
  - How useful is Optimum Write Size?

WinHEC  
2008



## The Effort and Progress are Ongoing

- Continue enabling SSD technology in Windows platform; Boot, Index, ...
- Working with IHVs and OEMs to enrich Windows hardware logo program for better quality
- Define SSD endurance specification together with industry

Microsoft  
WinHEC  
2008

## Call To Action

- Test your SSD with Windows 7 beta and optimize your device with new Windows 7 features
- Run Windows 7 logo evaluation kit (available soon) to ensure your device compliance
- Provide us with your feedback on future device optimization need
- Participate in industry standard work (T13, T10, JEDEC, SSDA, ...) and support SSD endurance, data retention, and device statistic standardizations

Microsoft  
WinHEC  
2008

## Additional Resources

### • Web Resources

- Trim command specifications:  
[http://t13.org/Documents/UploadedDocuments/docs2008/e07154r6-Data\\_Set\\_Management\\_Proposal\\_for\\_ATA-ACS2.doc](http://t13.org/Documents/UploadedDocuments/docs2008/e07154r6-Data_Set_Management_Proposal_for_ATA-ACS2.doc)
- ATA8-ACS: [http://t13.org/Standards/Default.aspx?DocumentType= 3&DocumentStage=2](http://t13.org/Standards/Default.aspx?DocumentType=3&DocumentStage=2)
- Microsoft Research Papers: Design Trade-offs for SSD Performance  
<http://research.microsoft.com/users/vijayanp/papers/ssd-usenix08.pdf>
- Guidelines for Designing Flash-Based Ultra Low Cost PCs for Windows XP:  
<http://download.microsoft.com/download/2/0/a/20ac945c-34d0-4a60-8245-f80e80fe954f/GuidelinesForDesigningULCPCForWindowsXP-MSFT%20.pdf>

### • Related Sessions

- “New Developments in the Storage Platform” - COR-T521
- “System Integrated Flash Storage” - COR-T559
- “Solid State Storage in the Server and Data Center Environments” - ENT-C628
- “Designing Flash-based Netbooks for Windows” - MBL-T549

### • Contact info

- Forum: Storage Platform -ATA  
<http://forums.microsoft.com/MSDN/ShowForum.aspx?ForumID=1837&SiteID=1>
- Email: [ssddiscu@microsoft.com](mailto:ssddiscu@microsoft.com)



**Microsoft**<sup>®</sup>  
Your potential. Our passion.<sup>™</sup>

© 2008 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries. The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.