

4kB Data Sector Update

IDEMA 4kB Technical Committee

Storage Developers Conference, September 2008

Presented by

Mike Fitzpatrick, Fujitsu Computer Products of America

- ❑ This presentation is being brought to you by the IDEMA 4K Technical Committee
- ❑ The purpose of this presentation is to make everyone aware of the transition of the physical sector size of Hard Disk Drives (HDD) from 512 bytes to 4K bytes
- ❑ Who might be affected:
 - ❑ Disk drive vendors - obviously
 - ❑ Operating System (OS) suppliers
 - ❑ Storage Subsystem developers
 - ❑ System developers
 - ❑ Application developers
 - ❑ Etc., etc.
- ❑ Additional Info – www.BigSector.org & www.IDEMA.org

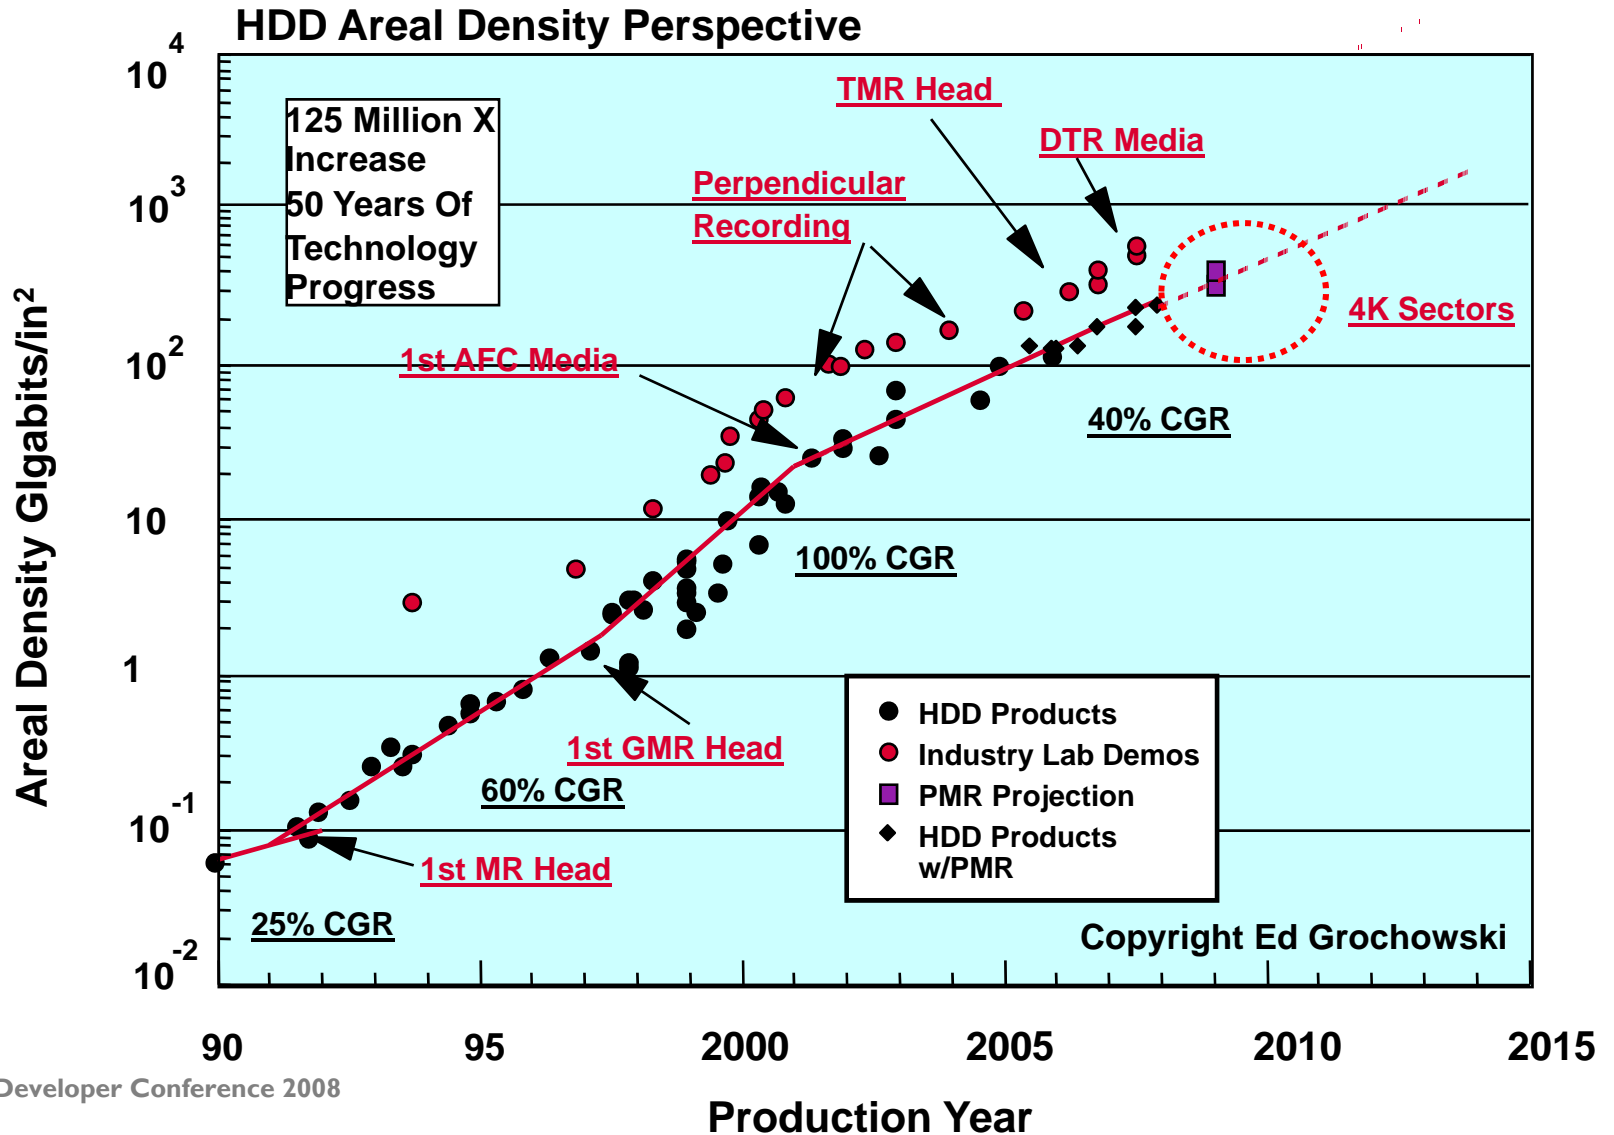
- ❑ IDEMA is the not-for-profit trade association serving the over \$60 billion data storage industry's growth worldwide with international trade shows, technical conferences, standards programs, dinner programs, and industry assessments. IDEMA members are comprised of companies in both hard drive and solid state drive manufacturing, including their diverse and robust supply chain.
- ❑ Who Are IDEMA Members? -- IDEMA corporate members are the companies that comprise the over \$60 billion data storage industry (HDD and NAND Flash) and its robust, dynamic supply chain, including the designers and manufacturers of hard drives, heads, media, substrate, process and test equipment, materials and other parts used in data storage devices, including solid state drives, and other alternative forms of storage. All employees of IDEMA corporate member companies are IDEMA members and are entitled to full member benefits including discounts on conferences, symposia DISKCON tradeshows and all other IDEMA activities offered worldwide. Complete details are available at www.idema.org

Brief History to 4K Sectors

- ❑ As far back as 2000, the HDD industry recognized the need to make changes to the data format in order to continue areal density growth.
- ❑ The Long Data Sector (LDS) Committee was formed under the guidance of IDEMA.
- ❑ In 2003, the committee requested Microsoft to support up to a 4K sector format in their next OS.
- ❑ Windows Vista and Longhorn both have some support for both emulated and non-emulated 4K sector hard disk drives.
- ❑ Today we are here to inform you of these changes coming and make you aware of the consequences.

Why 4K Sectors?

Areal Density Growth

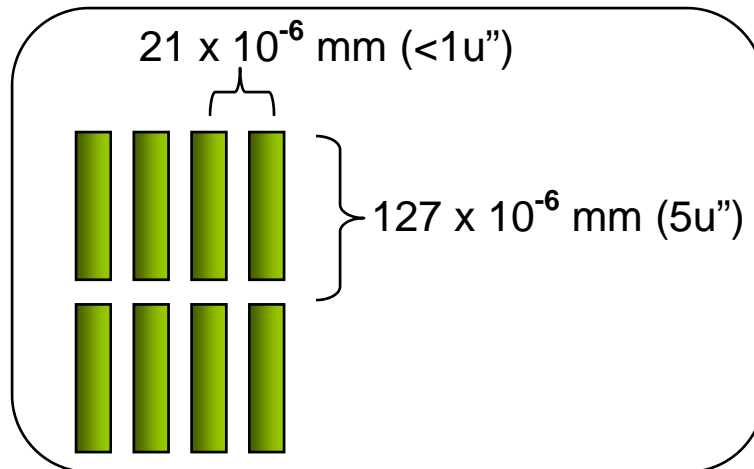


Areal Density Examined

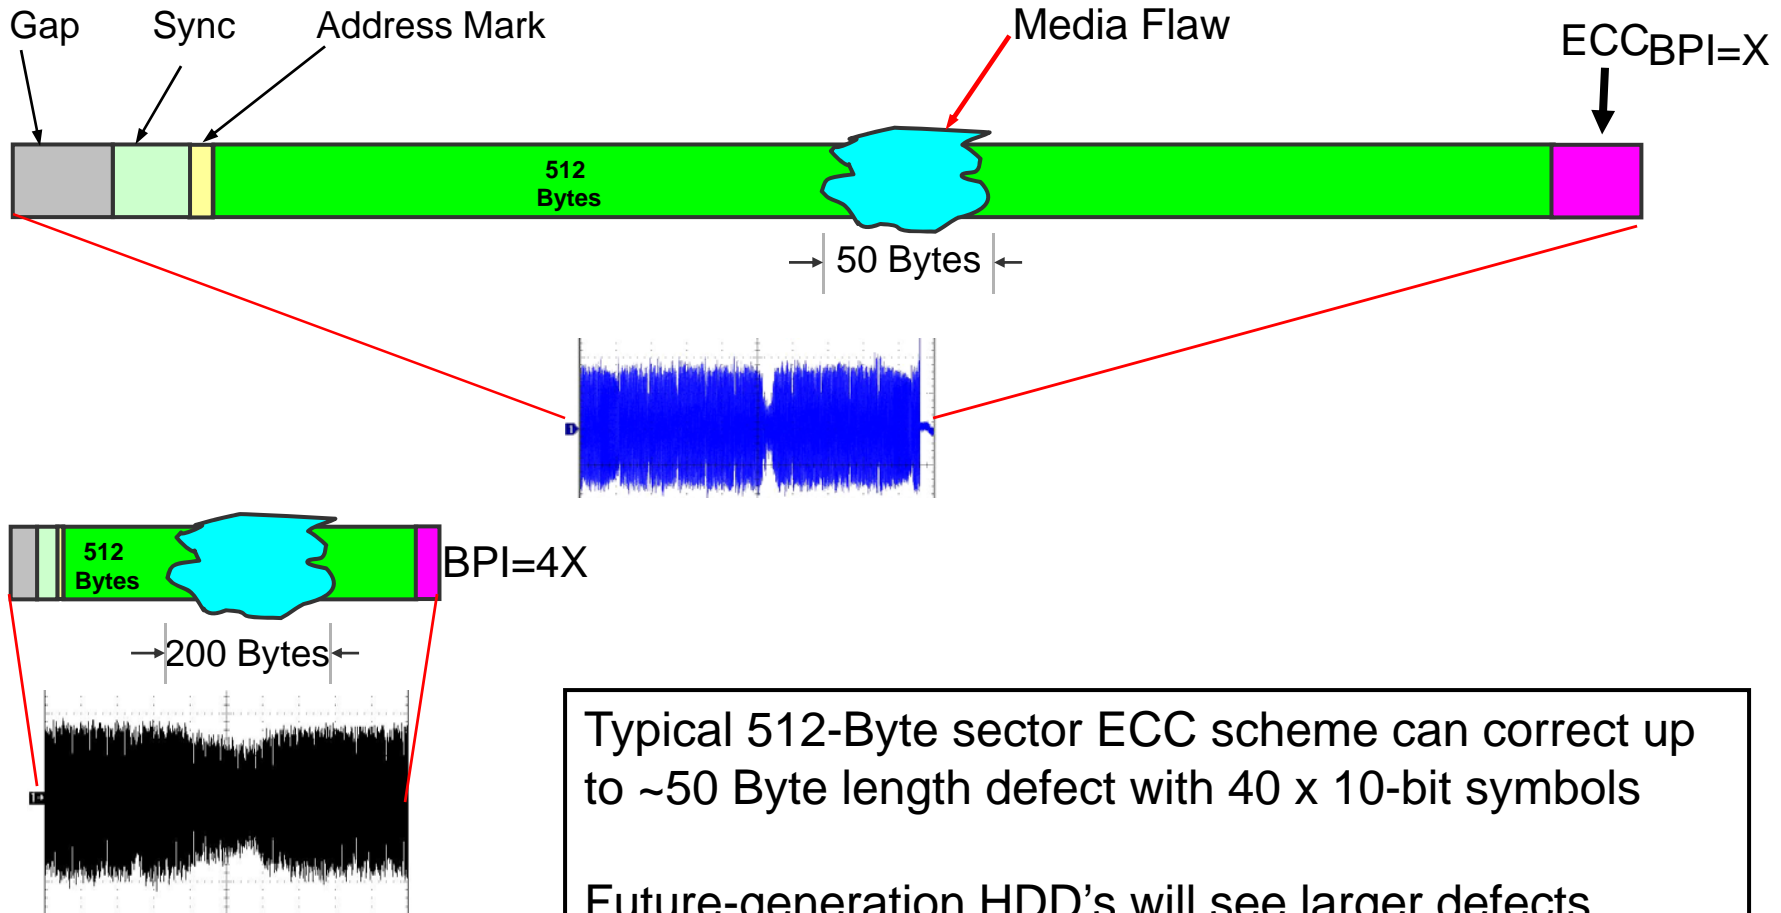
Example -

Areal Density of Mobile HDD: 240 Gb/in²

- Track density = 200k TPI, Track Pitch = 127nm (127×10^{-6} mm)
- Bit Density = 1200k BPI, Bit Length = 21 nm (21×10^{-6} mm)



Defect Density is Increasing

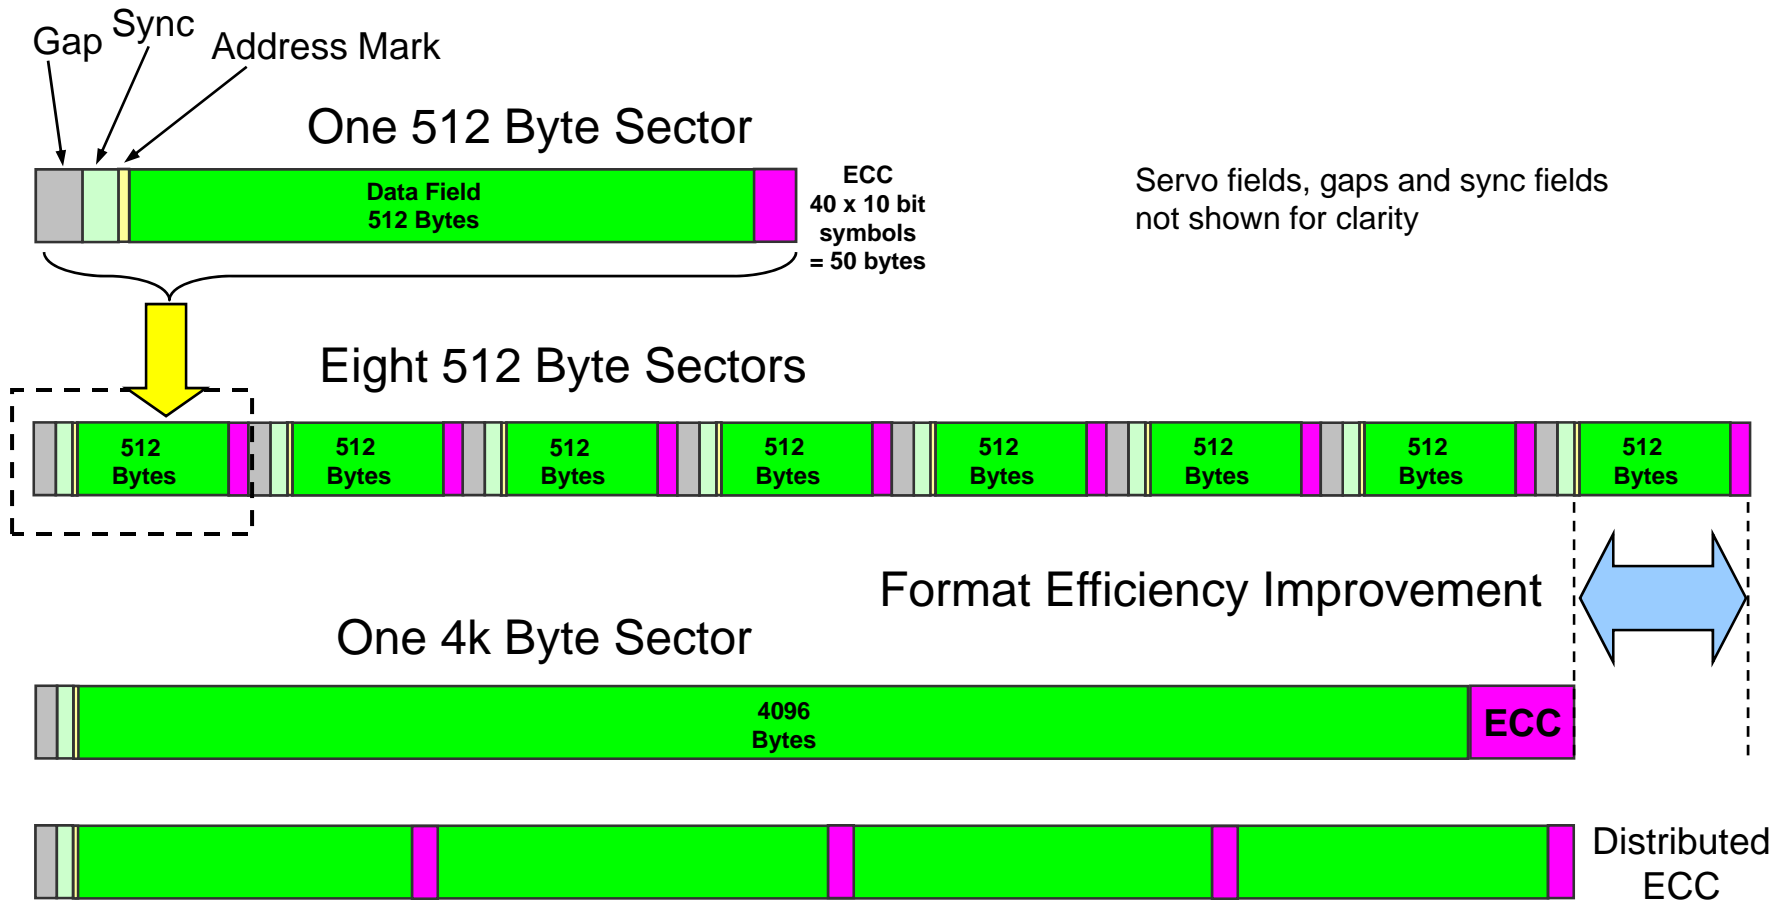


Typical 512-Byte sector ECC scheme can correct up to ~50 Byte length defect with 40 x 10-bit symbols

Future-generation HDD's will see larger defects

Must improve ECC scheme to maintain UER

Format Efficiency with Long Block



- Format Efficiency improves by 6-13% with 4kB sector (depends on 512B sector layout, and disk size)
- Gains can be used to reduce BPI or TPI and improve yield

The Impact of 4K Sectors

- ❑ New disk drive introductions from disk drive suppliers will ship with 4KB sectors on disk starting CYQ1 2011
- ❑ Implications:
 - ❑ Requires the entire industry food chain to prepare for the transition including OEMs, ODMs, Software apps, Partition tool providers, etc.
 - ❑ OEMs requiring 512B Native products will need to purchase legacy capacity points/drives
 - ❑ Decision does not preclude earlier 4KB product introductions

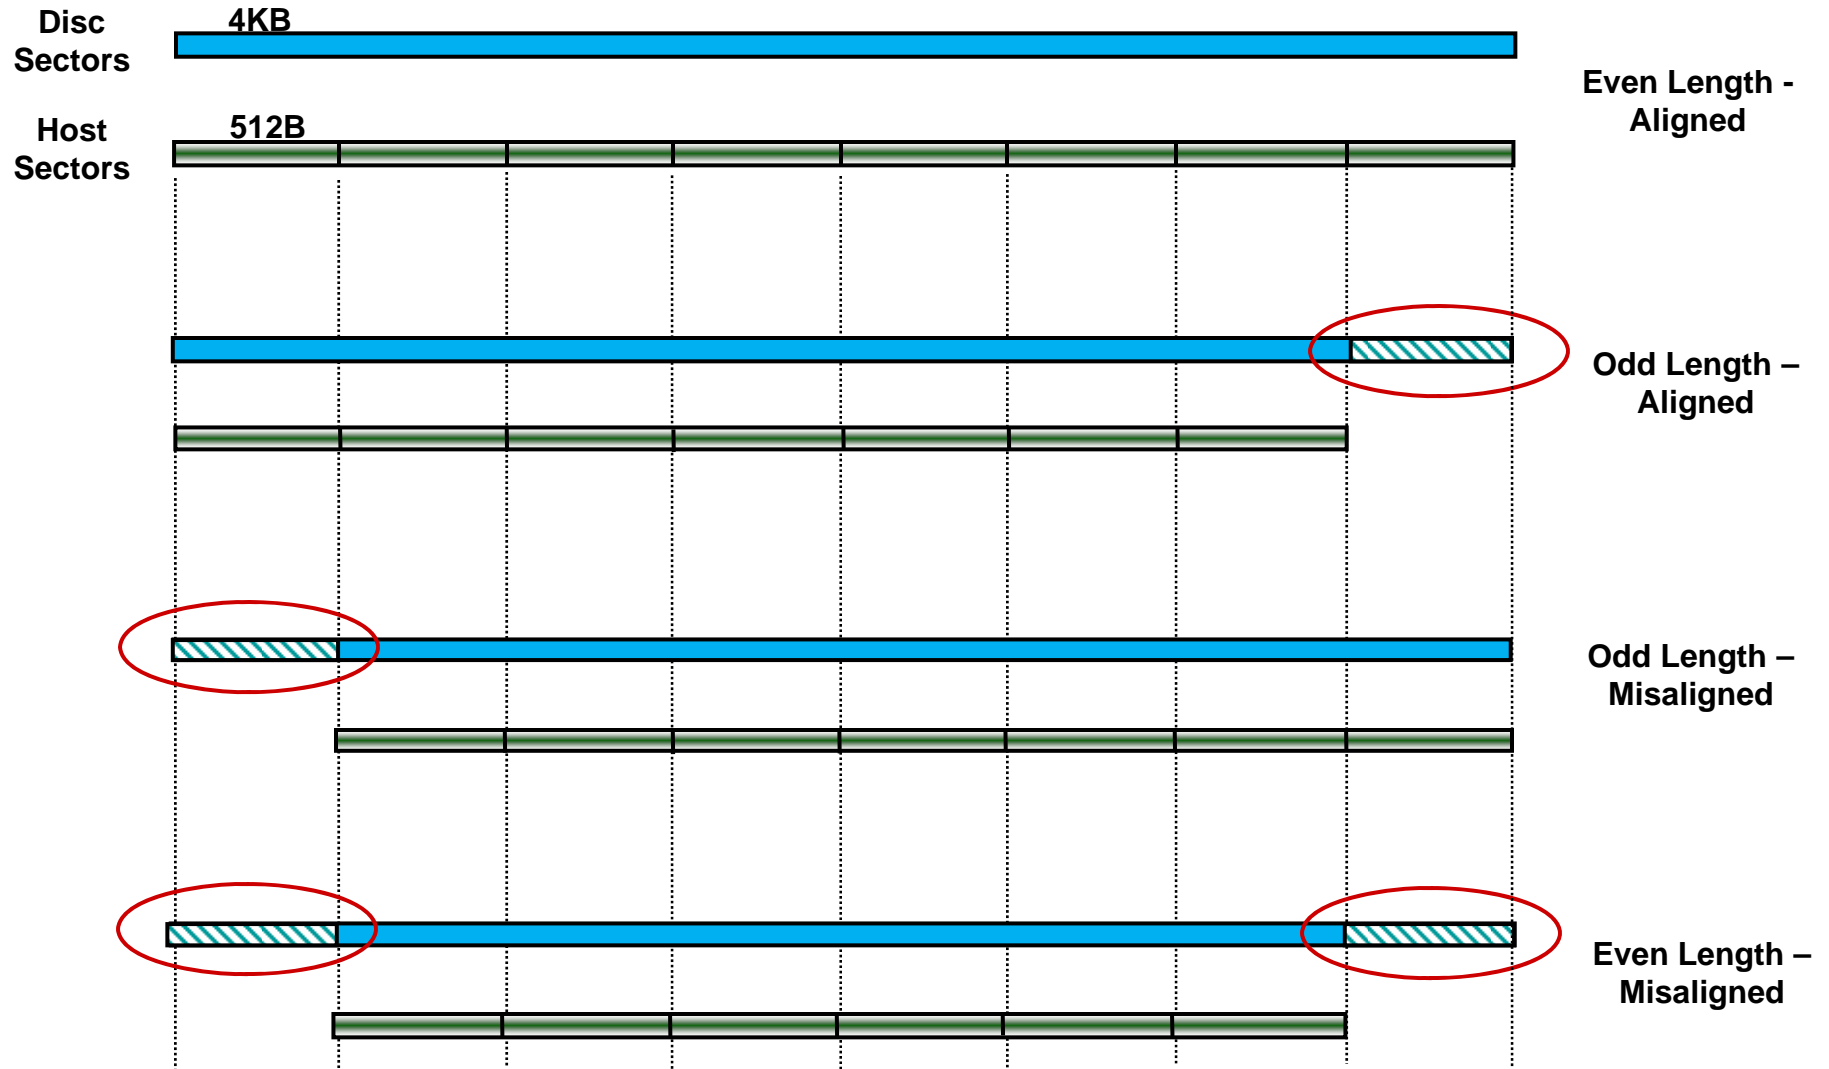
Key Points (cont)

- ❑ Disk Drive suppliers are planning to ship 4KB sector drives with 512B Emulation mode until full OS support for 4KB Native is available
- ❑ Implications:
 - ❑ Disk drive suppliers are planning to ship 512B Emulation until 4KB Native is fully supported by all OS and Apps (i.e. beyond 2012)
 - ❑ Enterprise SAS/FC drives may ship as 4KB Native to customers where unique support for 4KB sectors exists
 - ❑ A one-time quick format option to change between 512B Emulation and 4KB Native will be made available
 - ❑ No reformat options for SATA Enterprise applications are possible—i.e. Nearline
 - ❑ SATA Nearline must accept 512B Emulation or continue purchasing legacy 512B Native drives until a 4KB Native OS solution is available

Key Points (cont)

- ❑ All Disk Drive suppliers agree to a single LBA alignment default in 512B Emulation—one for SATA, another for SCSI
- ❑ Implications:
 - ❑ A single alignment standard should minimize the potential pitfalls due to Imaging Install processes
 - ❑ Provides best opportunity to ensure proper alignment even if Vista tools are bypassed by ODMs, Fortune 500 IT Departments, etc.
 - ❑ Once new images are created for large sector drives, they should be transportable to all disk drive suppliers without modification
 - ❑ Unique customers may require a different alignment—handled on exception basis
 - ❑ For SATA, one alignment option—alignment 1—provides the best performance for legacy, non large sector aware systems.
 - ❑ For SCSI = TBD

Alignment



Effects of misaligned transfers

- ❑ Writes starting on a 4K Byte boundary with transfer length a multiple of 4K Bytes require NO read-modify-write (RMW).
- ❑ Writes starting on a 4K Byte boundary with transfer length NOT a multiple of 4K Bytes require RMW on last LBA.
- ❑ Writes NOT starting on a 4K Byte boundary with transfer length a multiple of 4K Bytes require RMW on first and last LBA.
- ❑ Writes NOT starting on a 4K Byte boundary with transfer length NOT a multiple of 4K Bytes requires a RMW on first and (probably) last LBA.

- I. Propose to ANSI INCITS T10 and ANSI INCITS T13 Committee proposals to implement a correctible error option to inform hosts of misaligned transfers
 - Data transferred correctly but host receives a direct indication that a performance impact is occurring due to misaligned transfers (Read-Modify-Write events)
 - Proposal owner: WD

2. Propose an ANSI INCITS T10 standard for Format Device from 512B host transfers to 4KB host transfers
 - Note: Leave flexible host sector size for future host sector size increase
 - Proposal owner: Seagate

Host Interface


Drive Physical Sector

512

4K

	512 Native	512 Emulation	4K Native
512	512 Bytes on media -----	-----	-----
4K	-----	4K Bytes on media	4K Bytes on media

Eight - 512 Byte Sectors



One - 4k Byte Sector

- ❑ Document LBA1-03 (an IDEMA Standard)
 - ❑ This document defines the algorithm used by all HDD vendors for determining the number of LBAs a disk drive shall have based on the reported capacity of the drive.
 - ❑ This document has been expanded to include:
 - ❑ SATA/SAS Disk drives
 - ❑ SATA/SAS Large Data Sector (4K sector) drives
 - ❑ SAS disk drives whose sector is formatted with the T10 PI (Protection Information), a.k.a. DIF (Data Integrity Format).
 - ❑ The way in which a drive notifies the host of its LBA count is defined in standard documents referenced in section 3.0.
- ❑ Document available at www.IDEMA.org

Summary

- 4kB Logical Sector size impacts OS, BIOS and driver code
- BIOS will need to detect and utilize 4kB devices
- Some applications that bypass the OS File System need to be modified
- Emulation potentially impacts performance due to Read/Modify/Write
 - Will be necessary to support legacy applications
 - Not acceptable to some customers
- On Enterprise HDD's –
 - RAID controllers and other initiator devices need modifications
 - 512B Emulation may not be appropriate
- Impact of 4kB HDD's on Operating Systems and applications is being assessed by the IDEMA 4kB Block Working Group
 - All major HDD makers are represented
- Microsoft, Oracle and other software vendors are involved

- ❑ www.BigSector.org
 - ❑ 4KB Sector Transition
 - ❑ Hard Disk Drive Long Data Sector White Paper
 - ❑ Large Sectors Seagate Technology
 - ❑ Big Sector Background
 - ❑ Corporation Opportunities and Challenges for Higher Density Recording
 - ❑ Transition Requirements: Lessons from Serial ATA
 - ❑ Larger Sector Sizes: A Drive Vendor's Perspective
 - ❑ 4K-Block Format Efficiency and SNR Gain
 - ❑ Large Blocks for Reliability

4kB is coming !

Acknowledgements

Thanks to all who have contributed material towards various presentations and particularly those that have participated in the IDEMA 4kB (Long Data Sector) working committee and meetings.

Special Thanks go to:

Steve McCarthy – Consultant
Tom Pratt – Dell
Frank Widjaja – Dell
David James – Fujitsu
Mike James – Fujitsu
Paul Bailey – Hewlett Packard
Walter Gaspard – Hewlett Packard
Sami Waheed – Hewlett Packard
Sally Castillo – Hitachi
Dan Colegrove – Hitachi
Martin Hassner – Hitachi
Steve Livaccari – IBM
Dave Verburg – IBM
Debbie Lee – IDEMA
Dr. Ed Grochowski – IDEMA
Trudy Gressely – IDEMA
Joel Weiss -- IDEMA
Ana Whitlock -- IDEMA

Harley Burger – LSI
Bob Sheffield – LSI
Bret Weber – LSI
Zining Wu -- Marvell
Calvin Chen – Microsoft
John Gardill – Microsoft
Bob Griswold – Microsoft
Nathan Obr – Microsoft
Frank Shu – Microsoft
Jim Williams -- Oracle
Joseph Chen – Samsung
Yawshing Tang -- Samsung
David.A.Anderson – Seagate
Marty Czekalski – Seagate
Mary Dunn – Seagate
Marc Noblitt – Seagate
Patrick Hery – Toshiba
Curtis Stevens – WD